NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**INTERDICTING A FORCE DEPLOYMENT:
TWO-SIDED OPTIMIZATION OF ASSET SELECTION,
LIFT SCHEDULING,
AND MULTI-COMMODITY LOAD PLANNING**

by

Peter M. Koprowski

March 2005

Thesis Advisor:	Gerald Brown
Second Readers:	Glenn Lintz
	Matthew Carlyle

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE				*Form Approved OMB No. 0704-0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** March 2005	**3. REPORT TYPE AND DATES COVERED** Master's Thesis		
4. TITLE AND SUBTITLE: Title (Mix case letters) **Interdicting a Force Deployment:** **Two-Sided Optimization of Asset Selection, Lift Scheduling, and Multi-Commodity Load Planning**				**5. FUNDING NUMBERS**
6. AUTHOR(S) LT Peter M. Koprowski				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000				**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) Joint Warfare Analysis Center, Dahlgren VA.				**10. SPONSORING/MONITORING AGENCY REPORT NUMBER**
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited				**12b. DISTRIBUTION CODE** A
13. ABSTRACT (maximum 200 words) A military deployment is visible and vulnerable. But, deployments are currently planned assuming they can be completed with surprise, or defended from any threat. JFAST, the current deployment planning and visualization tool of choice, uses heuristics of unknown reliability that yield deployment plans of unknown quality, and ignores vulnerability. We introduce LIFTER, an integer-linear program (ILP) that optimizes a time-phased force deployment (TPFDD) by day, by asset cycle, and by TPFDD line (individual shipment from an origin to a destination), and ATTACKER, also an ILP, representing a smart enemy's resource-limited interdictions to maximally disrupt LIFTER's subsequently re-optimized TPFDD plan. LIFTER activates transport assets from an allocation list, and yields a complete logistic plan that minimizes disruption represented by penalties for early, tardy, late, or dropped shipments, and for under-utilization of asset capacity. We use LIFTER to qualitatively assess JFAST heuristic plans. We also link both ILPs in a decomposition-based search for the best deployment plan around the worst-case interdiction, given that the actions of deployer and interdictor are transparent to both parties. We explain how JFAST could be embellished with its own version of ATTACKER. A key discovery here is a gauge of the value of intelligence, deception, and secrecy.				
14. SUBJECT TERMS TPFDD, Time-Phased Force Deployment, JFAST, optimization, interdiction, two-sided optimization, value of intelligence				**15. NUMBER OF PAGES** 76
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified		**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

INTERDICTING A FORCE DEPLOYMENT: TWO-SIDED OPTIMIZATION OF ASSET SELECTION, LIFT SCHEDULING, AND MULTI-COMMODITY LOAD PLANNING

Peter M. Koprowski
Lieutenant, United States Navy
B.S.,United States Naval Academy, 1997

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
March 2005

Author: Peter M. Koprowski

Approved by: Gerald G. Brown
Thesis Advisor

Glenn J. Lintz, CDR, SC, USN
Second Reader

Matthew Carlyle
Second Reader

James N. Eagle
Chairman, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

A military deployment is visible and vulnerable. But, deployments are currently planned assuming they can be completed with surprise, or defended from any threat. JFAST, the current deployment planning and visualization tool of choice, uses heuristics of unknown reliability that yield deployment plans of unknown quality, and ignores vulnerability. We introduce LIFTER, an integer-linear program (ILP) that optimizes a time-phased force deployment (TPFDD) by day, by asset cycle, and by TPFDD line (individual shipment from an origin to a destination), and ATTACKER, also an ILP, representing a smart enemy's resource-limited interdictions to maximally disrupt LIFTER's subsequently re-optimized TPFDD plan. LIFTER activates transport assets from an allocation list, and yields a complete logistic plan that minimizes disruption represented by penalties for early, tardy, late, or dropped shipments, and for under-utilization of asset capacity. We use LIFTER to qualitatively assess JFAST heuristic plans. We also link both ILPs in a decomposition-based search for the best deployment plan around the worst-case interdiction, given that the actions of deployer and interdictor are transparent to both parties. We explain how JFAST could be embellished with its own version of ATTACKER. A key discovery here is a gauge of the value of intelligence, deception, and secrecy.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

The following is a Time-Phased Force Deployment (TPFDD) lexicon. Some of these terms do not appear in this document, but all of them are essential to decipher documentation of seminal TPFDD planning methods.

ACL	Allowable Cabin Load: The maximum payload that can be carried on an airlift mission. It may be limited by the maximum takeoff gross weight, maximum landing weight, or by the maximum zero fuel weight.
AK	AK: air via strategic (AMC, AMC-contract) aircraft
ALD	Available-to-Load Date: A day, relative to C-day in a time-phased force and deployment data, that unit and nonunit equipment and forces can begin loading on an aircraft or ship at the port of embarkation [Jones, 2005]
AMC	Air Mobility Command: Military command responsible for providing airlift, air refueling, special air mission, and aeromedical evacuation for U.S. forces.
B8 Format	A specific format for TPFDD data produced by JOPES.
Breakbulk_mt	Ship's capacity for breakbulk equipment measured in mTons
Bulk	Bulksize: Cargo that is within the usable dimensions of a 436L pallet (104'' x 84'' x 96'') and within the height and width requirements established by the cargo envelope of a particular model aircraft. A sub-identifier after sTons to further describe the cargo associated with a single Line requirement
BULK_st	Aircraft capacity for bulk sized equipment measured in sTonss
CCC	Cargo Category Code: Identifier that describes the type, extent, and containerization of cargo
C-Day	Commencement Day: The day on which deployment operations begin
CONOPS	Concept of Operations
Container_mt	Ship's capacity for containerized equipment measured in mTons
CONUS	Continental United States
CRAF	Civil Reserve Air Fleet
EAD	Earliest Arrival Date: A day, relative to C-day, that is specified by a planner as the earliest date when a unit, a resupply shipment, or replacement personnel can be accepted at a port of debarkation during a deployment. [Jones, 2005].
ELD	Earliest-Load-Date: The day relative to C-Day that a lift asset is available for loading [Jones, 2005].
ELIST	Enhanced Logistics Intratheater Support Tool

GEO	Geographic Identifier: A four digit code that identifies a specific location
JOPES	Joint Operations Planning and Execution System
JSCP	Joint Strategic Capabilities Plan
LAD	Latest Arrival Date: A day, relative to C-day, that is specified by a planner as the latest date when a unit, a resupply shipment, or replacement personnel can arrive and complete unloading at the port of debarkation and support the concept of operations [Jones, 2005].
Line	A TPFDD requirement for transportation
LOLO	Lift-on Lift-off
Lolo_sqft	Ship's capacity for lift on lift off equipment measured in square feet
LRWC	Long Range Wide-body Cargo: This term is used to describe cargo aircraft utilized from the CRAF fleet.
LRWP	Long Range Wide-bodied Passenger: This term is used to describe passenger aircraft utilized from the CRAF fleet.
MAX_ACL_lbs	The maximum allowable cabin load measured in pounds for cargo carriers and passengers for passenger carriers.
MAX_FUEL_lbs	The maximum amount of fuel measured in pounds that can be help by an aircraft.
MAX_TOGW_lbs	The maximum takeoff gross weight. The operating weight plus the fuel weight plus the cargo weight measured in pounds.
MIN_ACL_lbs	A percentage of the ACL that must be filled by cargo available at the APOE to commit an aircraft to a mission
MS2POD	Mode Source to POD: Two character identifier that relays the mode of transportation for a line requirement to its Port of Debarkation
MSC	Military Sealift Command: Military command responsible for providing ocean transportation of equipment, fuel, supplies and ammunition to sustain U.S. forces worldwide during peacetime and in war for as long as operational requirements dictate.
mTons	Measured Tons: a unit of volume equal to 40 cubic feet
OPER_WT_lbs	Operational weight: The operating weight of the aircraft in pounds. The operating weight is the weight of the aircraft less fuel and cargo.
OPLAN	Operational Plan
Out	Outsize: Cargo which exceeds the dimension of oversize (1,090" x 117" x 105") and requires the use of a C-5 or a C-17. A sub-identifier after sTons to further describe the cargo associated with a single Line requirement
OUT_st	Aircraft capacity for out sized equipment measured in sTons
Over	Oversize: Cargo exceeding the usable dimensions of a 436L pallet loaded to the design height of 96" but is equal to or less that 1,090" in length, 117" in width, and 105" in height. This cargo is transportable on the C-5, C-17, C-141, C-130, and to a limited extent the KC-10. A sub-identifier after sTons

	to further describe the cargo associated with a single Line requirement
OVER_st	Aircraft capacity for over-sized equipment measured in sTons
PAX	Passengers
PID	Plan Identification: A command-unique four-digit number followed by a suffix indicating the JSCP year for which the plan is written, e.g., "2220-95." In JOPES data base, a five-digit number representing the command unique four-digit identifier, followed by a one character alphabetic suffix indicating the OPLAN option, or a one-digit number numeric value indicating the JSCP year for which the plan is written.
POD	Port of Debarkation: The geographic point at which cargo or personnel are discharged. May be a seaport (SPOD) or aerial port (APOD) of debarkation. For unit requirements, it may or may not coincide with the destination.
POE	Port of Embarkation: The geographic point in a routing scheme from which cargo or personnel depart. May be a seaport (SPOE) or aerial port (APOE) from which personnel and equipment flow to port of debarkation. For unit and nonunit requirements, it may or may not coincide with the origin.
POL	Petroleum, Oil and Lubricants
RDD	Required Delivery Date: A date, relative to C-day, when a unit must arrive at its destination and complete offloading to properly support the concept of operations [Jones, 2005].
RLN	Requirement Line Number: TPFDD requirement description
RORO	Roll-on Roll-off
Roro_sqft	Ship's capacity for roll-on roll-off equipment measured in square feet
SDDC	Surface Deployment and Distribution Center: Military command responsible fro providing global surface deployment command & control and distribution operations to meet National Security objectives in peace and war
SE	SE: Sea via MSC ship (common user strategic sealift)
sqft	Square Feet: Two dimensional description of cargo associated with a Line Requirement
sTons	Short Tons: a unit measure of weight equivalent to 2000 pounds
TPFDD	Time-Phased Force Deployment Data
WLTAE	Warfighting and Logistics Assessment Environment

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

This work would not have been possible without the help of several dedicated people. Specifically Gerald Brown, Distinguished Professor of Operations Research, whose tireless efforts have been a motivating force to get this project done. Mr. Phil Boyer, USTRANSCOM, showed patience and generously offered continued support that proved essential to this work. Thanks for never getting frustrated with the questions. Colonel Saverio Manago, without you, we never would have even gotten off the ground. Finally I thank my wife for her never-ending support over the last mile.

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

Even though our military logistic deployments are vulnerable to interdiction, we plan such deployments without any formal consideration of interdiction threats and currently assume we can achieve and maintain military supremacy sufficient to secure our movements of personnel and materiel. Planning a deployment assuming it will be invulnerable to interdiction may be current, accepted practice, but this is neither realistic, nor safe.

Surprisingly, there is scant literature on force deployment planning with optimization. Two publications worth noting are a 1999 Naval Postgraduate School masters thesis by LT. William Brown, and a 2002 paper by Professors David Morton (University of Texas at Austin), Kevin Wood, and Javier Salmeron (both of the Naval Postgraduate School). Brown introduces deterministic re-optimization to respond to simulated emergent changes during a deployment, such as denial of debarkation port access. Morton, Wood and Salmeron plan a sealift deployment while hedging against potential disruptions by a random attack. Each publication highlights that current logistics planning tools do not anticipate disruptions, and that this can be a key deficiency. Additionally, each paper concentrates on sealift and focuses on major adversarial military actions in a two-sided theater engagement. We are concerned with airlift as well as sealift, and the attacks we anticipate include limited insurgent actions. We also assume an intelligent enemy will anticipate deployment activities and attack optimally with intent to maximize disruption, rather than attack at random.

We use an "training" time-phased force deployment data TPFDD designed by planners at Fort Eustis for a scenario involving approximately half a US Army division to deploy from its continental US (CONUS) origins to debarkation points in Azerbaijan and Georgia, near the Caspian Sea.

We assume any debarkation airport or seaport might be vulnerable to an insurgent attack, and that such an attack would cause immediate disruption, essentially closing the debarkation point while security is restored, and that subsequent debarkation throughput would take a few days to fully recover. We do not conjecture any particular weapon or

tactics on the part of the insurgents, reckoning that any breach of security at a debarkation point would inflict some disruption.

To demonstrate how the attacker's actions can be modeled, we endow the insurgents with the will to mount between one and three such attacks over the deployment planning horizon, but allow no more than one attack in any five-day period.

We show how to plan a deployment to anticipate and minimize the worst-case consequences of any interdiction. We develop an integer linear program, LIFTER, for optimizing of a deployment, and compare our optimal results with those of Joint Flow and Analysis System for Transportation (JFAST), evidently the current planning tool of choice for deployment planning. JFAST plans with what appears to be a fast, rule-based, myopic heuristic, but is not a true optimization. LIFTER employs integer linear programming (ILP) to plan an optimal time-phased force deployment by transport mode (air and sea); with each TPFDD line (shipment) a homogeneous commodity and cargo type (TPFDD level three resolution). Time fidelity is daily. The objective is to minimize the penalties for early, tardy, or late deliveries, or dropping lines all together, weighted by the "importance" (or military worth) of each line.

We then present ATTACKER, another integer linear program that plans resource-limited attacks on our deployment intended to maximize disruption, assuming that we stay committed to our deployment plan even after the attacks begin. Given any optimal (i.e., minimum disruption penalty) plan from LIFTER, ATTACKER plans a maximally-disrupting set of interdictions.

Next, we assume transparency between the deployer and the interdictor --- each opponent can see what the other is planning to do, and adjust counter-plans accordingly. The Naval Postgraduate School has produced a large body of research in the last few years applying optimization to exploit open-source data in red-teaming exercises to analyze vulnerability and defense of critical infrastructure. Their results confirm that the statement, "Using public sources openly … it is possible to gather at least 80% of information about the enemy," excerpted from a captured Al Qaeda training manual, is a conservative estimate. Here we apply these results to deployment planning, assuming that such a major logistic deployment cannot be hidden.

We find the best deployment plan that responds optimally to the worst that an interdictor can do.

JFAST delivers all but 7% of the general breakbulk (mTons) cargo in our uninterdicted TPFDD deployment. LIFTER delivers all of the cargo on time. JFAST and LIFTER use lift assets at approximately the same rate. JFAST plans 247 missions to satisfy 100% of the air requirements while LIFTER plans 261. LIFTER plans twice as many C-5 missions and one-and-a-half times as many C-17 missions. For sea delivery JFAST schedules 9 ships to move all but 7% of the general breakbulk cargo, with an average ship capacity utilization of 85%. LIFTER schedules nine ships to deliver 100% of the sea requirements with an average ship capacity utilization of 57%.

There are distinguishing differences between JFAST and LIFTER.

LIFTER delivers an objective assessment of the quality of its plan. That is, you know with certainty how much better any undiscovered plan can be. JFAST plans are of completely unknown quality.

JFAST is an identity simulation: it features accurate fidelity, and represents logistic flows from seminal origins to each end unit in theater. LIFTER only models the strategic lift from embarkation to debarkation points, and approximates time fidelity to daily resolution.

JFAST violates cargo loading restrictions by mixing non-compatible cargo types in common compartments, but LIFTER does not.

JFAST can animate its planned deployment. But, neither we nor other interested researchers have been able to find documentation explaining how JFAST assembles a plan. When JFAST suggests inappropriate or blatantly infeasible plans, we are hard-pressed to guess why.

The only way we have found to glean cargo manifests from JFAST is to manually click on the moving, animated icon of each lifting asset.

We cannot reckon how JFAST chooses ships to activate: we have seen JFAST ignore available ships and over-use others, or make very late debarkation deliveries.

The ILP ATTACKER, endowed with three attacks for the deployment horizon, finds the most-disruptive set of attacks. ATTACKER achieves its highest rate of success by attacking Poti on day C+20 and C+27 and Baku Bina on C+35. This optimal attack plan forces LIFTER to drop 1% of the passenger requirements and 1% of the material sTon requirements. All sea requirements are delivered around this attack plan. The Subsequent analysis yielded similar results when the maximum number of attacks was decreased from three to two and then one. The delays are modest, but the throughput penalties are high.

Our two-sided optimization offers insight into the values of secrecy, deception, and intelligence. We view our first LIFTER plan as a "secret deployment," with no anticipation at all of any interdiction, and our first ATTACKER interdiction as a "surprise attack," made on an unsuspecting deployment. We then compare these two extreme cases with a completely "transparent" pair of deployment and interdiction plans.

If the deployer anticipates interdiction, and could keep the deployment a complete secret, the difference between the military values of the "transparent" case and the "secret" one indicates how much this would be worth. If the interdictor anticipates the deployment, and wonders how valuable a deception would be that completely hides the threat of interdiction from the deployer, the difference between the military worth of the "surprise" case and the "transparent" offers this. These values of secrecy become the goals of intelligence efforts by the respective adversaries.

We find that if we could deploy with total secrecy, we are marginally penalized for a few tardy lines resulting in a value of 1.5 million "importance-day" units (i.e., our composite measure of the importance of each line multiplied by the days of tardiness). If the enemy interdicts the TPFDD by surprise with his optimal one attack plan, the impact is 520 million "importance-day" units. The difference between these two numbers represents the value of operating secretly, or the value of intelligence.

Given the speed of the JFAST planner and its features to exclude zones and prevent flow to designated regions, implementing a JFAST equivalent of ATTACKER should not be too difficult.

I. INTRODUCTION

A. PROBLEM STATEMENT

Even though a military logistic deployment is visible and vulnerable to interdictors, we currently plan such a deployment assuming we can complete it in secrecy or defend it from any threat, and we assume that we can achieve and maintain military supremacy sufficient to secure our movements of personnel and materiel. Planning a deployment assuming it will be invulnerable to interdiction may be current, accepted practice, but this is neither realistic nor safe.

We show how to plan a deployment to anticipate and minimize the worst-case consequences of any interdiction. Mobilization and deployment of large-scale military forces from their home stations to an operational theater is a large logistics problem. We illustrate our new planning technique using a strategic mobilization and deployment of several Army units from their home stations within the continental United States to an operational theater in Eastern Europe. The US military uses the Joint Operations Planning and Execution System (JOPES) to plan and conduct military operations. For deliberate planning, each theater combatant commander is responsible for preparing an approved concept of operations (CONOPS) from the Combined Joint Chiefs of Staff. Specific tasking for these CONOPS plans is levied through the Joint Strategic Capabilities Plan (JSCP). During the CONOPS development, a plan database is created that represents the forces necessary to conduct the operation. This database, referred to as the Time-Phased Force Deployment Data (TPFDD), explicitly details the timeline necessary to move forces to the theater in accordance with the concept of operations [Joint Forces Staff College Pub 1, 2000]. The TPFDD is tested for feasibility using various planning aids, including the Joint Flow and Analysis System for Transportation (JFAST) [Jones, 2005]. Deployment planners use JFAST to determine transport requirements, perform course-of-action analysis, and project delivery profiles of troops and equipment by air, land, and sea.

To the best of our knowledge, JFAST uses an undocumented local search heuristics of unknown reliability to suggest a deployment plan of unknown quality. We

develop a mathematical optimization of a deployment and compare our optimal results with those of JFAST.

We then model resource-limited attacks of our deployment aimed at maximizing disruption, assuming that we stay committed to our deployment plan even after the attacks begin.

Next, we assume transparency between ourselves and the interdictor --- each opponent can see what the other is doing, and can thus anticipate and respond to threats by re-planning the deployment or interdictions.

By comparing the results of optimal "surprise interdictions" on a static TPFDD deployment with the transparent results, we can gauge the value of secrecy and deception to either opponent. This illustrates a new, quantitative way to evaluate the "value of intelligence" in terms understood by both the deployer and the interdictor.

B. LITERATURE REVIEW

Mathematical optimization has played a key role in planning transportation for over a half century. Koopmans [1946] earned the 1975 Nobel Prize in economics for his now-ubiquitous linear programming transportation models. Surprisingly, there is scant literature on TPFDD planning with optimization.

Brown [1999] uses deterministic re-optimization to respond to simulated emergent changes during a deployment. Brown reveals a shortcoming in the Enhanced Logistics Intratheater Support Tool (ELIST), the logistics simulation used in conjunction with the warfare simulation THUNDER that interfaces with the Warfighting and Logistics Assessment Environment (WLTAE) [JHU APL 2005]. ELIST has no ability to reroute supplies to alternate debarkation points in the event of an attack denying access to a planned debarkation point. Brown optimally reschedules ship destinations and arrival times in response to such an attack.

Morton, Wood, and Salmeron [2002] plan a sealift deployment while hedging against potential disruptions to a logistics network caused by an attack. Attacks are assumed to follow a probability distribution developed through interpretation of

intelligence reporting. The model is a multi-stage stochastic-program, with recourse. This work also highlights the inability of current TPFDD planning tools, specifically JFAST, to re-optimize deployment plans given attacks denying access to selected facilities.

Brown [1999] and Morton, et al [2002] focus on sealift. We are concerned with both sealift and airlift. Their declared goals focus on major adversarial military actions in a two-sided theater engagement. Ours include limited insurgent attacks. We also assume an intelligent enemy will not attack following some probability distribution, but will instead attack with intent to maximize disruption.

C. TPFDD DESCRIPTION

1. Definition

The TPFDD is a key component of force planning in the Joint Operations Planning and Execution System (JOPES) [2000, Joint Force's Staff College Pub. 1, B-3]. JOPES is used by national and theater commanders to establish a level of command and control necessary for the planning and conduct of joint operations. JOPES includes all the joint operational planning policies, "procedures and reporting structures." JOPES is used to monitor, plan, and execute mobilization, deployment, employment, and sustainment of joint operations [2000, Joint Forces Staff College Pub 1, G-45]. JOPES helps build joint operations.

The TPFDD is:

> The JOPES database portion of an operation plan; it contains time-phased force data, non-unit related cargo and personnel data, and movement data for the operation plan including (a) in-place units, (b) units to be deployed to support the operation plan with a priority indicating the desired sequence for their arrival at the port of debarkation, (c) routing of forces to be deployed, (d) movement data associated with deploying forces, (e) estimates of non-unit related cargo and personnel movements to be conducted concurrently with the deployed forces, and (f) estimates of transportation requirements that must be fulfilled by common-user lift resources as well as those requirements that can be fulfilled by assigned or attached transportation assets. [2000, Joint Forces Staff College Pub 1, GL-78]

2. Terminology

TPFDD lexicon is acronym-rich and filled with military jargon crucial to understanding a deployment. A deployment moves personnel and equipment from the US to some foreign operating area in three phases: 1- mobilization, 2- strategic, and 3- theater. This requires transport assets and infrastructure. During the mobilization phase forces arrive at their origins and prepare to begin the deployment. The strategic phase of a deployment begins with forces moving from their origins to respective sea or air embarkation ports (SPOE or APOE), and concludes when these forces arrive at their sea or air debarkation ports, SPOD or APOD (see Figures 1 and 2). Movement of forces to their final destination occurs during the theater phase. The embarkation and debarkation points have the key infrastructure needed to support the deployment. The transport assets include heavy transport aircraft and cargo ships (see Figure 3).

A TPFDD comprises a list of shipment requirements, called “lines.” Each line requirement is either moved by air (denoted “AK”) or by sea (“SE”).

The Joint Strategic Capabilities Plan (JSCP) apportions resources, or assets, to combatant commanders based on current military capabilities [Joint Forces Staff College Pub 1, 2000].

A TPFDD represents the time-phasing of the deployment. Required personnel and equipment must arrive in theater synchronously. Lessons learned from Operations Restore Hope in Somalia highlight the disastrous results of equipment arriving at a debarkation point without the proper personnel in place to receive it [Allard 1995]. The Earliest Arrival Date (EAD), Ready to Load Date (RLD), and Available Load Date (ALD), are all critical days in the TPFDD (see Figure 4).

Figure 1. Seaport of Embarkation (SPOE). Cargo in the marshalling yard is positioned for loading on to ships. Each ship has storage areas, or compartments, that are each dedicated to one type of cargo load, including some subset of breakbulk; container; lift-on, lift-off (LOLO); roll-on, roll-off (RORO); and petroleum, oil and lubricants (POL). Some ships have their own cranes, while others depend on port material handling equipment. Compartment capacities for breakbulk and container loads are expressed in measured tons, while LOLO and RORO are in square feet.

Figure 2. Airport of Embarkation (APOE). A cargo chalk (i.e., items designated to load on a given aircraft) is a commingled load, so the capacity of the aircraft is expressed in terms of total load capacity (maximum load less fuel load for the longest flight leg, in short tons); bulk cargo; outsize cargo; and/or oversize cargo.

Figure 3. Clockwise from upper left C-5, C-17, C-141, RORO ship offloading a tank, Breakbulk ship (Cape Gibson), and Light Aboard Ship (LASH, similar to a civilian container vessel).

The "who and what" are reported in plain text in the TPFDD lines: The operation planners decide who should deploy to execute the mission, and then those designated forces take what they need to accomplish their mission.

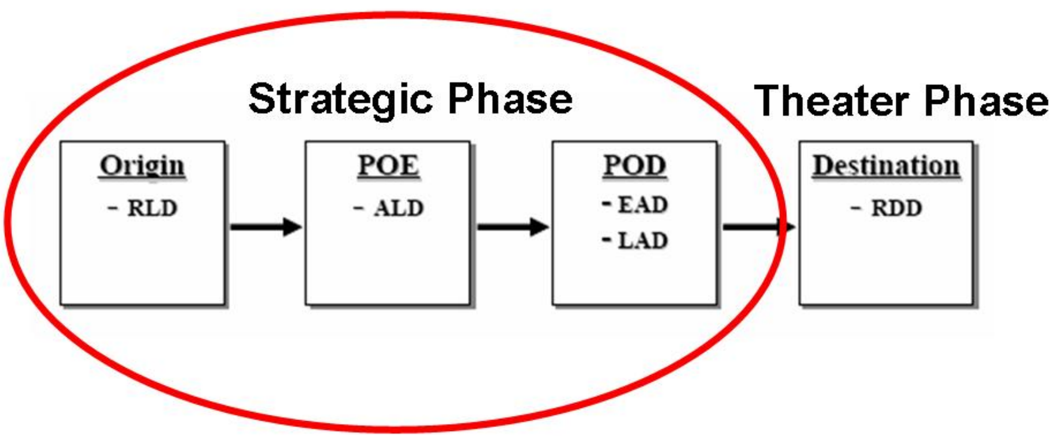

Figure 4. TPFDD timing acronyms govern the movement of personnel and materiel from origin to theater, and include Ready Load Date (RLD), Available Load Date (ALD), Earliest and Latest Arrival Date (EAD, and LAD), and Required Delivery Date (RDD). Our focus is on the POE-POD segment of planning

3. Dissection

A TPFDD times movements relative to a commencement day (C-Day) [2002, Joint Forces Staff College Pub 1, GL 78].

The TPFDD scenario we use here is comprised of 157 summary (TPFDD level two) line requirements. Each line represents personnel or equipment associated with a specific unit to be moved from a particular embarkation location to a particular debarkation location in the theater. Each line can be divided into three portions: (1) identifier and timing, (2) location and movement, and (3) associated cargo (by passengers or weight). Using the line called "FALHP" from our scenario, we can divide it into these three categories to understand the requirement as reported through JFAST.

PID	RLN	CCC	RLD	ALD	EAD	LAD	RDD
MW011	FALHP	PAX	25	26	28	32	45

Table 1. Identifier and timing acronyms used in the TPFDD. The Plan Identifier (PID) associates this line with the specific plan MW011. The Requirement Line Number (RLN) represents an index for the actual line requirement in the TPFDD. The cargo category code (CCC) reports the type of cargo the line represents. The timing portion shows the days the line is ready to be handled: Ready to Load Date (RLD) is the day the line is ready to begin loading at its Origin. Available Load Date (ALD) is the day when the line is ready for loading at the embarkation point. Earliest Arrival Date (EAD) and Latest Arrival Date (LAD) refer to the earliest and latest days that a line will be accepted at its debarkation point. Required Delivery Date is the day when the line is expected to be in place in theater to support current operations.

Table 1 illustrates the identifier and timing acronyms (column headings), and the associated entries for the line denoted "FALHP" in our training TPFDD. The identification portion of "FALHP" shows this line is part of a plan identifier (PID) with the name "MW011" and the associated cargo with cargo category code "PAX" or passengers (see Table 2). The timing portion reports all critical days, relative to C-Day, for "FALHP". For example the earliest arrival day (EAD) at its debarkation point is day C+28.

Cargo Category Codes		
Type (1st Character)	**Extent (2nd Character)**	**Containerization (3rd Character)**
A: Vehicles Non Self-Deploy	1: Outsize Unit Equipment	B: Can Containerize, 20FT container, 20 STons or less
B: Non Self-Deploy Aircraft	2: Oversized Unit Equipment	C: Can Containerize, 40FT container, 30 STons or less
J: Other Non Vehicle	3: Bulk Unit Equipment	D: Non Containerizable
M: Ammunition		
R: Vehicles Self-Deploy		

Table 2. Three-character Cargo Category Code descriptions for TPFDD lines. For instance, a helicopter is designated "B1D." In addition to these materiel codes, the code for passengers is "PAX."

Origin	Origin Name	MS2POE	POE	POE NAME	MS2POD	POD	POD Name
ESGM	FT DRUM	LH	YVGO	Wheeler Sack AAF	AK	AUTS	Baku-Bina

Table 3. Associated cargo, location, and movement fields for the TPFDD. Origin and Origin Name are field identifiers for the GEO code and text name of the starting position for a line requirement. MS2POE indicates the mode source to the port of embarkation (POE). POE and POE NAME are similar to the origin description providing the GEO code and text name of the POE. MS2POD indicates the mode source to the port of debarkation (POD). POD and PODNAME are the GEO code and text name of the POD.

The TPFDD line location segment (see Table 3) displays the physical location of "FAHLP" from starting location in the "Origin" filed, embarkation point in the "POE" filed, and debarkation point in the "POD" filed. This also shows the means of travel between these points. For example in the mode source to POE (MS2POE) field "LH" reveals that "FALHP" is using organic vehicles to get to the POE. Location data is expressed with a four-character geographic identifier (GEO) code and by a text description. The means of movement is expressed by a two-digit movement code (AK for air, SE for sea).

PAX	Bulk mTons	Bulk sTons	Over mTons	Over sTons	Out mTons	Out sTons	Cbbls pol
129	0	0	0	0	0	0	0

Table 4. TPFDD cargo quantities. A TPFDD line, such as this one for "FALHP", divides commodities into passengers, bulk-, over-, or out-sized cargo, and Petroleum, Oil and Lubrication products (POL). Cargo quantity is given in both measured tons (mTons) and short tons (sTons). Measured tons expresses cargo quantity for sealift, and represents a volume equivalent to 40 cubic feet. Short tons expresses cargo quantity for airlift of 2,000 pounds. Cbbls is hundreds of barrels.

Table 4 illustrates the associated cargo portion of the line "FALHP." We see that "FALHP" is comprised of 129 passengers.

A TPFDD line composed of these three components is called "level two," or summary level. This is the highest level of aggregation for a line requirement. There is a disaggregation of each level two line into one or more level three lines. Each level three line is composed of logistically-homogeneous cargo: these lines can actually be loaded onto transport assets, and thus there is sufficient cargo fidelity to determine a complete shipping plan. Even more detail is visible at higher levels, but we do not need this for our work. We use TPFDD level three for modeling movement to the theater. At level three, our scenario TPFDD has 639 lines.

D. JFAST (JOINT FLOW AND ANALYSIS SYSTEM FOR TRANSPORTATION)

The transportation planning tool JFAST (Joint Flow and Analysis System for Transportation) has been developed for the US Transportation Command (USTRANSCOM) [Jones, 2005] to provide a planner with a graphical user interface that illustrates the entire planned deployment. JFAST supports planning in both the strategic and theater phases of a deployment. Following JOPES planning guidance, the combatant commander will consolidate plan inputs and submit the complete plan for review and feasibility testing. As part of that review USTRANSCOM will analyze the strategic sea and air transportation feasibility [2000, Joint Forces Staff College Pub 1, 4-15]. JFAST

streamlines plan preparation by forecasting and illustrating for the commander what USTRANSCOM will do to validate and approve any submitted TPFDD.

JOPES produces and exports a TPFDD in “B8 format,” [McKinzie and Barnes, 2004] and this is the only format JFAST accepts. Once a TPFDD is imported it can be viewed and edited in every detail. A planner can also build a TPFDD from scratch using JFAST and then export that plan in B8 format for interoperability with other JOPES systems.

In addition to line requirements, there are fields for each lift asset, origin-embarkation point, destination-debarkation point, and the C-Day timeline. Airlift and sealift modes each have a separate interface for managing asset attributes. Lift assets are assigned to a plan by either using the Joint Strategic Capabilities Plan apportionment resident in JFAST or by manual entry.

POE-POD information comes from Air Mobility Command (AMC) for airports and the Surface Deployment and Distribution Command (SDDC) for seaports. Military planners coordinate with AMC and SDDC representatives to agree on embarkation and debarkation point planning factors. Once a deployment plan substantially reflects the concept of operations, it can be executed and visualized in motion. The JFAST planner can run the deployment in its entirety, or run any of three separate planning spans (ORIGIN to POE, air POE to POD, or sea POE to POD).

For planning in the strategic phase, JFAST decides how to load and move airlift and sealift cargo and passengers from each origin to debarkation point. JFAST documentation gives few details about how this movement is planned. McKinzie and Barnes critically analyze JFAST and report that

> All JFAST schedules use only a simple greedy heuristic to generate solutions. JFAST assesses cargo and PAX based on priorities and schedules the two from highest priority down. No process is considered that would optimize the scenario.

This evaluation is echoed by Morton, Wood and Salmeron [2002] who state that JFAST, “cannot optimize (or re-optimize) a schedule with respect to an objective such as minimize cost.”

JFAST plans and schedules both air and sea deployments. For the air portion JFAST allows the user to select one of three models. Each model produces varying levels of fidelity and insight into the deployment.

> The airflow estimator is a capacitated flow model that distributes airlift capability across competing movement requirements, by day, with the objective of minimizing lateness as compared to the latest arrival date of the requirement (2005, JFAST 8.0 Handbook, 3-3B1).

The "airflow estimator" produces the least fidelity and is recommended for a very large TPFDD. Conversely the "full air scheduler" produces the most fidelity for the deployment.

> The full air scheduler converts the air fleet into individual aircraft, assigning notional tail numbers to each. The number of aircraft in combination with an aircraft type utilization rate is used to generate a flying hour pool, by aircraft type, for each day. During scheduling, all aircraft are evaluated for every mission to determine the aircraft that delivers the most cargo within the desired time window. It is also important to note that the full scheduler does not use average load data as published in the Joint Strategic Capabilities Plan (JSCP) mobility annex. Instead, the full scheduler performs routing analysis to determine the critical leg for each mission. From the critical leg, a fuel plan is estimated and the allowable cabin load (ACL) determined on a mission-by-mission basis. The routing function's objective is to minimize total mission time. (2005, JFAST 8.0 Handbook, 3-3B1).

The last air model is the "quick scheduler." The quick scheduler closely mimics the full air scheduler except a flying time estimator is used instead of dynamic routing [Jones, 2005]. An important difference is that the quick scheduler does not show fuel plan management. Without this visibility, it is impossible to view the available cabin load per aircraft.

The sea model uses one heuristic, optimizing scheduler [Jones, 2005] to produce the deployment plan. Insight to the deployment planned is provided via the sealift rainbow chart [Jones, 2005]. The rainbow chart graphically displays the itineraries for the ships participating in the deployment. Additionally the rainbow chart allows the user to view cargo manifests for a participating vessel.

Once JFAST completes the planning scenario, logistics performance statistics are provided through a series of pre-formatted reports or planner-defined ad-hoc reports. JFAST uses maps and simulated-time animation to illustrate a deployment. This visualization allows the planner to stop time motion and click on any asset to discover its manifest and intentions (see Figure 5). Unfortunately this is the only way to achieve in-transit visibility of lines being transported by air assets.

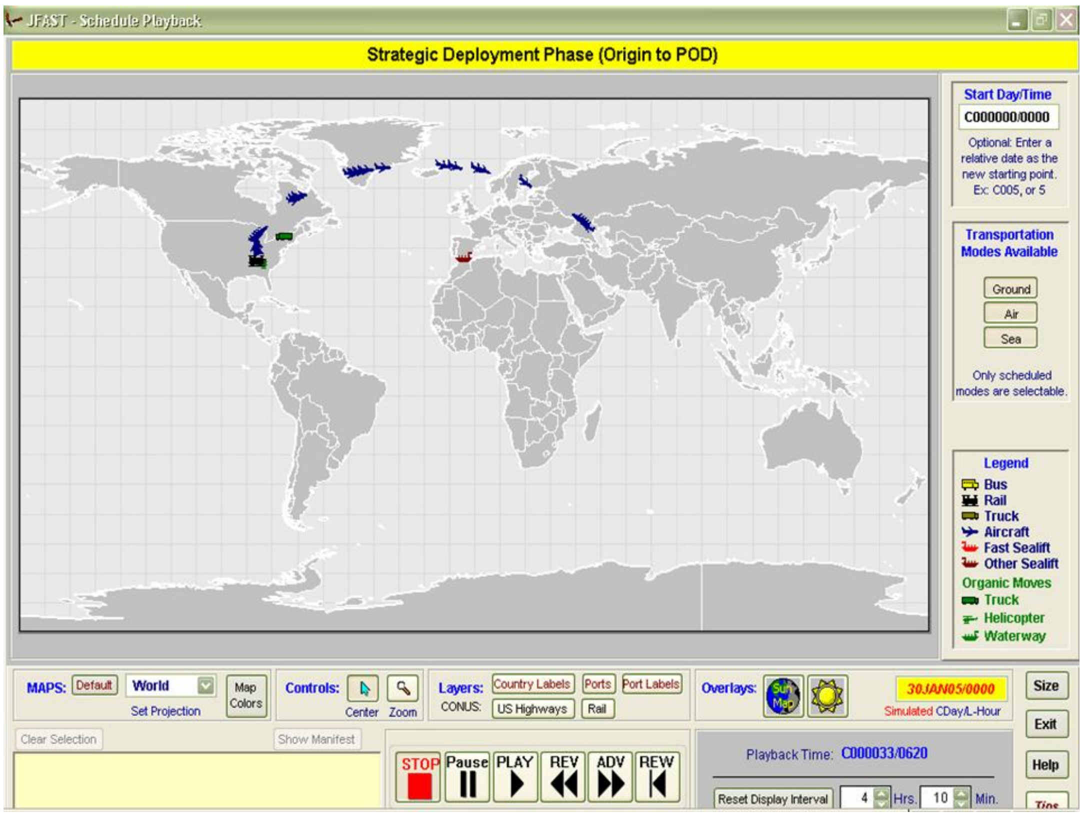

Figure 5. JFAST deployment playback screen shot. Aircraft cycles from the US to the Azerbaijan theater are animated in simulated time. There is also a ship shown transiting the Strait of Gibraltar. A mouse-click can stop the animation, and drill down to the cargo manifest of any asset in transit. However, there is no means to extract this level three data in a format suitable for analysis.

E. TPFDD SCENARIO

We use an unclassified "training" TPFDD designed by planners at Fort Eustis for a mobility warrant officers' course that calls for about half of a US Army division to deploy from its continental US (CONUS) origins to debarkation points in Azerbaijan and Georgia, near the Caspian Sea.

Personnel and equipment from four CONUS-based Army forts mobilize and move to their associated air and sea embarkation points for deployment (see Figures 6 and 7). The summary total volumes of passengers and cargo moved are shown in Table 5.

	PAX	STons	MTons
Air	8,726	3,937	0
Sea	0	0	133,260
Totals	**8,726**	**3,937**	**133,260**

Table 5. Summary of total cargo to be lifted to the theater. Cargo is reported in passengers (PAX), short tons (sTons), and measurement tons (mTons).

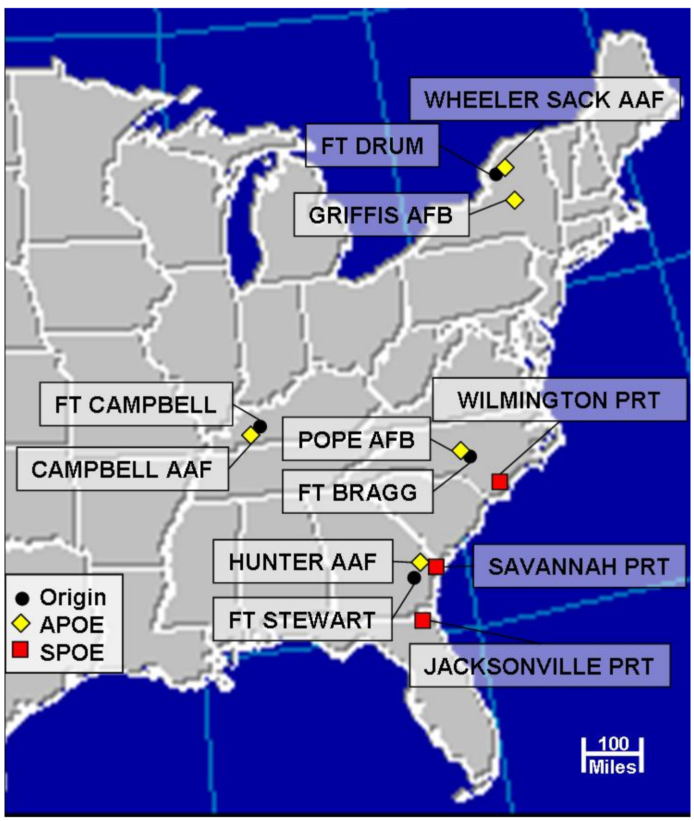

Figure 6. Origin locations and embarkation points.

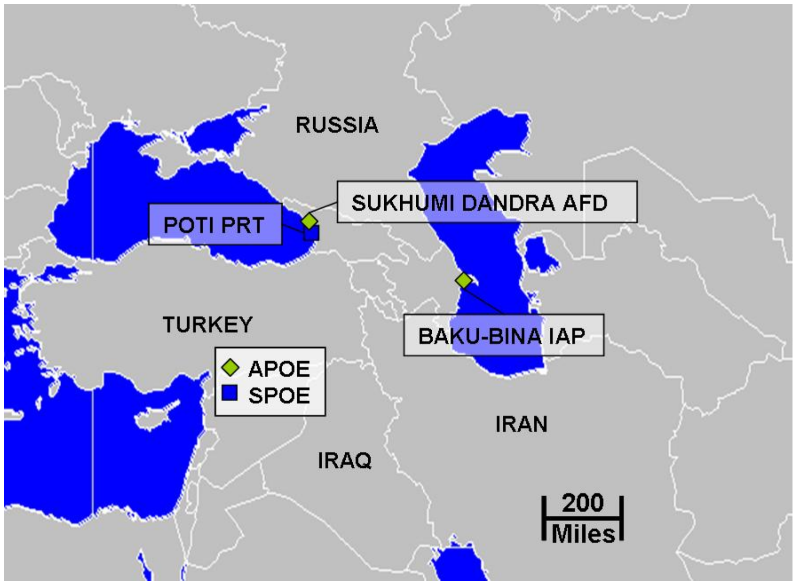

Figure 7. Debarkation points in Georgia and Azerbaijan.

F. ATTACKS

We assume any debarkation point might be vulnerable to an insurgent attack, and that such an attack would cause immediate disruption, essentially closing the debarkation point while security is restored, and that subsequent POD throughput would take a few days to fully recover. We do not conjecture any particular weapon or tactics on the part of the insurgents, reckoning that any breach of security at a debarkation point would inflict some disruption. For example, a kinetic attack with vehicle borne improvised explosive devices or short-range rockets would suffice. Alternately, enemy sympathizers might dissuade workers from reporting to POD jobs.

We also make a key assumption that distinguishes our work from any prior TPFDD analysis: ***if the US begins moving personnel and materiel by major sealift and airlift to conduct operations in a forward area, this will not be done by surprise.*** There are many recent examples where escalating diplomatic activity and economic sanctions have broadcast our intentions prior to military deployment. Given such prior warning, it is not difficult for an enemy to figure out the likely debarkation points.

Bennett [1997] suggests that an enemy in this situation has no response but to seek asymmetric means to throw off the US and buy as much time as possible in hopes to

persuade US and international public opinion. He presents a scenario where forces on the verge of war with the US engage in terrorist attacks against airports and seaports, seizing initiative and engaging the US before they are able to fight. Clarke [2005] forecasts the future of the war on terror to feature the continued use of low-tech terrorist means to achieve maximum public impact. Such insurgent measures are far easier to mount in a foreign country than in US territory.

We model such attacks to completely stop throughput at an APOD or SPOD for a period of time, after which material and personnel handling capabilities incrementally regenerate to normal levels over time (see Figure 8). To demonstrate how the attacker's actions can be modeled, we endow the insurgents with the will to mount a few such attacks over the deployment planning horizon, but allow no more than one attack in any five-day period. Because we represent the attacker's actions by a general integer linear program, the fidelity of our attacker model is limited only by our imagination.

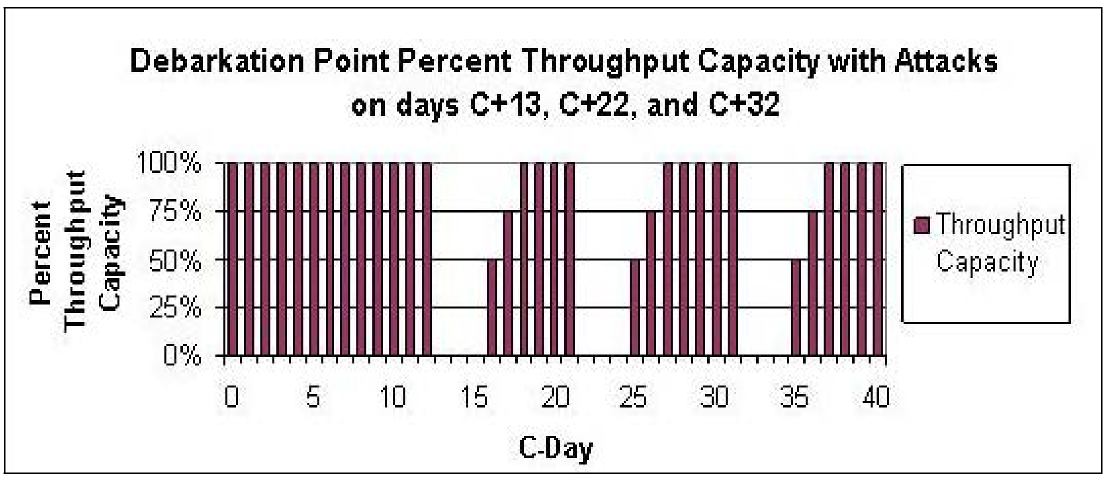

Figure 8. Throughput capacity at a debarkation point that is attacked three times over the deployment planning horizon. This figure depicts the effects of three attacks spaced no closer to each other than five days, on "C" days 6, 11, and 25. Throughput capacity regenerates to 50 percent three days after an attack, 75 percent a day later, and is back to 100 percent after five days. Such interdiction effect profiles are easy to generalize to suit intelligence estimates of potential enemy courses of action.

THIS PAGE INTENTIONALLY LEFT BLANK

II. OPTIMIZING FORCE DEPLOYMENT TRANSPORT

A. TPFDD

The size of each TPFDD line in our training example varies from a fraction of a ton to several hundred tons (see Figure 9). This is peculiar for a logistics planning model, and defies our experience with hundreds of case studies and examples in the military and private sectors. We decided not to follow standard practice rectifying this situation, and leave the data intact. This large diversity of shipment sizes does not bode well for an optimization model. Larger deployments would exacerbate this difficulty.

Using JFAST TPFDD editor we export level three TPFDD lines to a spreadsheet. Each of these level three requirements is categorized into one of nine cargo types based on its cargo category code (see Table 6). The spreadsheet is extracted to a comma-separated value file for direct use by our algebraic modeling language, General Algebraic Modeling System (GAMS) [Brooke, et. Al., 1998].

We endow our TPFDD with an additional attribute for modeling: we concoct a relative importance for each line in the TPFDD. This importance is calculated using the line attributes: EAD, mode source to POD, cargo category code, sTon weight, and POD. Earliest arrival dates for the TPFDD are C+22, C+25, C+28, or C+33. A line with an EAD of C+22 is given a weight of four, an EAD of C+25 receives a weight of three and so on. This places higher importance on cargo in the TPFDD that is desired earlier.

A TPFDD designated to travel via air (mode source to POD of AK) is given a weight of three and a line designated to travel via sea is given a weight of one. This places higher importance on cargo traveling by air.

The weighting of personnel arriving in theater is a two-step evaluation. TPFDD lines moving passengers are given a weight of two. This places a priority on personnel arriving in theater. However, if a TPFDD passenger line size is between 26 and 33 personnel it is given a weight of three. Based on the TPFDD, the design of this operation is very aviation centric. Aviation maintenance units consist of 26 to 33 personnel per the

TPFDD. Therefore the additional weighting of these lines places higher importance on aviation maintenance personnel arriving in theater.

TPFDD lines debarktion at Sukhmi Drandra airport are given a weight of three, Baku Bina a weight of two and Poti a weight of one. The debarkation point weighting continues adding higher importance to air TPFDD lines.

Finally, the weight (sTons) of each TPFDD line is multiplied by the importance factors a line acquires. The resulting equation takes the form:

$$\text{importance}_l = \text{sTon}_l \left(\text{weight}_{\text{EAD}_l} * \text{weight}_{\text{MS2POE}_l} * \text{weight}_{\text{PAX}_l} * \text{weight}_{\text{POE}_l}\right) \qquad \forall l$$

This line importance score is our surrogate for logistic military value, and we assume that the deployer and interdictor both view each line to have this value.

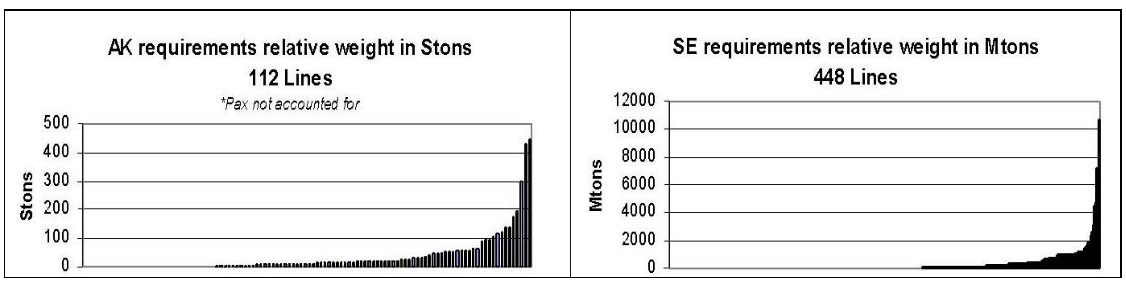

Figure 9. Weights of TPFDD lines vary from a few hundred pounds to thousands of tons. While this wide variation may be a necessity for military planning, it complicates modeling of such plans. Most decision support systems for logistics planning encourage aggregation of like items, but that is not the practice here. Vertical bars on the x-axis represent individual TPFDD lines. Lines close to the x and y axis intersection are so small they are invisible on this graphic.

Air Cargo Categories	Total Cargo per Category
Bulk sized (sTons)	2,440.9
Over sized (sTons)	13,569.8
Out sized (sTons)	19,595.9
Sea Cargo Categories	
Breakbulk (mTons)	7,632.0
Container General (mTons)	165.0
Container Ammo (mTons)	5,507.0
Container Vehicles (mTons)	2,461.0
Vehicles (sqft)	425,970.0
Non Self-deployable Aircraft (sqft)	96,363.0

Table 6. Cargo to be moved is classified into nine cargo categories

As stated earlier the training TPFDD is comprised of 639 TPFDD level three lines. After review of the data it is possible to reduce this number for faster model solve times. By sorting the lines according to embarkation point, available-to-load date, earliest and latest arrival date, required delivery date, and cargo handling characteristics, then combining lines that match in all fields, we can reduce the TPFDD to 32 "*TPFDD level 2.5*" lines. This aggregation is nearly equivalent to level three, and is easier to experiment with.

B. MATHEMATICAL FORMULATION OF LIFTER, AN INTEGER LINEAR PROGRAM TO OPTIMIZE TPFDD DELIVERIES

Index Use [~cardinality]

$m \in M$	transport mode (i.e., air, sea) [~2]
$l \in L$	TPFDD shipment line [~1,000]
$a \in A$	transport asset [~200]
$d \in D$	day (alias δ) [~80]
$n \in N$	node for transport (alias s, t) [~20]
$s(l) \in N$	source node for line l
$t(l) \in N$	destination node for line l
$b \in B$	type of cargo load burden (e.g., weight, cube, footprint, containers) [~20]
$c \in C$	type of load compartment (e.g., oversize, breakbulk, RORO) [~10]
$y \in Y$	node daily throughput capacity (e.g., passengers, airport standard

tons, seaport metric tons) [~10]

Given Data [units]

$importance_l$	value of line l [value-units]
$\theta_e, \theta_t, \theta_l$	penalties for early, tardy, or late delivery [1/day]
θ_d	penalty for dropping volume [fraction]
θ_u	penalty for empty asset capacity [value-units/b-unit]
$burden_{l,b}$	cargo burden b of shipment line l [b-units]
$\overline{capacity}_{a,c}$	maximum capacity of asset a compartments of type c [c-units]
$stow_{b,c}$	cargo burden b that can be loaded into compartment type c [b-units/c-unit]
$\overline{lifts}_a$	maximum number of lifts for asset a [cardinality]
$\underline{day}_l, \overline{day}_l$	time window for delivery of line l to node $t(l)$ [days]
$\overline{throughput}_{n,y,d}$	throughput capacity y at node n during day d [y-units/day]
$handling_{n,b,y}$	node n throughput capacity y required by cargo burden b [y-units/b-unit]
$days_{a,s,t}$	time for asset a to transit from node s to node t [days]

Derived data [units]

$efficiency_{l,b,a,c} \equiv burden_{l,b} / (stow_{b,c} \overline{burden}_{a,c})$ fraction asset a compartment type c constitituted by load l cargo burden b [fraction].

Decision Variables [units]

$ACTIVATE_a$	=1 if asset a used [binary]
$LIFT_{a,d,s,t}$	=1 if asset a to depart node s on day d, bound for node t [binary]
$EMPTY_{a,c,d,s,t}$	fraction of asset a capacity type c unused by $LIFT_{a,d,s,t}$ [fraction]
$LOAD_{l,a,d}$	fraction of line l loaded on transporter a on day d [fraction]
$DROP_l$	fraction of line l never delivered [fraction]
$TARDY_l$	fraction of line l delivered later than earliest delivery [days]
$EARLY_l, LATE_l$	fraction of line l delivered out of time window [days]

Formulation

$$
\begin{array}{llll}
MIN & \sum_{l} importance_l \left(\theta_e EARLY_l + \theta_t TARDY_l + \theta_l LATE_l + \theta_d DROP_l\right) & & \text{(L0)} \\
 & + \sum_{a,c,d,s,t} \theta_u EMPTY_{a,c,d,s,t} & & \\
s.t. & \sum_{a,d} LOAD_{l,a,d} + DROP_l = 1 & \forall l & \text{(L1)} \\
 & \sum_{\substack{l|s=s(l)\\ \wedge t=t(l)}} \left[\sum_{b} efficiency_{l,b,a,c}\right] LOAD_{l,a,d} + EMPTY_{a,c,d,s,t} = LIFT_{a,d,s,t} & \forall a,c,d,s,t & \text{(L2s)} \\
 & \sum_{\substack{l|s=s(l)\\ \wedge t=t(l)}} \left[\sum_{c,b} efficiency_{l,b,a,c}\right] LOAD_{l,a,d} \le LIFT_{a,d,s,t} & \forall a,d,s,t & \text{(L2c)} \\
 & \sum_{\substack{l|n=s(l),\\ a}} handling_{n,b,y} burden_{l,b} LOAD_{l,a,d} & & \\
 & + \sum_{\substack{l|n=t(l),\\ a,\\ \delta=d-days_{a,s(l),t(l)}}} handling_{n,b,y} burden_{l,b} LOAD_{l,a,\delta} \le throughput_{n,y,d} & \forall n,y,d & \text{(L3)} \\
 & \sum_{s,t,d \le \delta \le d+2days_{a,s,t}-1} LIFT_{a,\delta,s,t} \le 1 & \forall a,d & \text{(L4)} \\
 & \sum_{d,s,t} LIFT_{a,d,s,t} \le \overline{lifts}_a ACTIVATE_a & \forall a & \text{(L5)} \\
 & EARLY_l = \sum_{\substack{a,d|\\ \underline{day}_l-(d+days_{a,s(l),t(l)})>0}} \left[\underline{day}_l - (d+days_{a,s(l),t(l)})\right] LOAD_{l,a,d} & \forall l & \text{(L6e)} \\
 & TARDY_l = \sum_{\substack{a,d|\\ (d+days_{a,s(l),t(l)})-\underline{day}_l>0}} \left[(d+days_{a,s(l),t(l)}) - \underline{day}_l\right] LOAD_{l,a,d} & \forall l & \text{(L6t)} \\
 & LATE_l = \sum_{\substack{a,d|\\ (d+days_{a,s(l),t(l)})-\overline{day}_l>0}} \left[(d+days_{a,s(l),t(l)}) - \overline{day}_l\right] LOAD_{l,a,d} & \forall l & \text{(L6l)} \\
 & ACTIVATE_a \in \{0,1\} & \forall a & \text{(L7)} \\
 & LIFT_{a,d,s,t} \in \{0,1\} & \forall a,d,s,t & \\
 & EMPTY_{a,c,d,s,t} \ge 0 & \forall a,c,d,s,t & \\
 & LOAD_{l,a,d} \in [0,1] \text{ or } \{0,1\} & \forall l,a,d & \\
 & DROP_l \in [0,1],\ \ EARLY \ge 0_l,\ \ TARDY_l \ge 0,\ \ LATE_l \ge 0 & \forall l &
\end{array}
$$

1. LIFTER Formulation Discussion

LIFTER takes a time-phased force deployment data (TPFDD) and plans the deployment by transport mode, with each TPFDD line (shipment) a homogeneous commodity and cargo type (TPFDD level three resolution). Time fidelity is daily: Each transport asset departure is scheduled to a nearest day. A large TPFDD line may have to be split across asset departures, while smaller lines are shipped as unit loads. Delivery of each TPFDD line is completed as early as possible within its time window, or a variable accounts for any fraction of line volume delivered early (prior to time window), tardy (after start of time window), or late (after time window), or what fraction of line volume is never delivered at all (dropped): For any such bad outcome, the objective (L0) evaluates a penalty proportionate to the importance of the entire line and the maximal fraction-days of deviation. (L0) also penalizes any unused (empty) capacity on any activated asset.

Each constraint (L1) assures that a TPFDD line is either completely delivered by transport assets at some time, or accounts for any volume dropped (never delivered). Each constraint (L2) limits the loaded cargo burden for each transport departure to the compartment capacity of that transport asset. (L2s) is used for an asset with segregated storage of burden types to make sure that each type of burden capacity is honored (such a constraint is used for a sealift asset). (L2c) is used for an asset with commingled burden types to make sure that the total burden capacity is honored (such a constraint is used for an airlift asset).

For each node, day, and burden type, a constraint (L3) limits burden loaded and/or unloaded. (Note: This throughput may be vulnerable to and influenced by interdiction actions not shown here.)

Given some transport asset departure on some day from some source to some destination, a packing constraint (L4) precludes any subsequent departure of that asset for any location until it can have been recovered to that location.

A constraint (L5) limits the maximum number of lifts for an activated asset during the entire planning horizon.

Each line has a delivery time window expressed as an earliest and latest delivery day. Within a time window, an earlier delivery is preferred. If any fraction of a line is

delivered before its time window, a constraint (L6e) accounts for the total earliness of this in fraction-days. A constraint (L6t) accounts for any tardy delivery of a line within its time window, but not at the start of this window. If any part of a line is delivered after the time window, a constraint (L6l) accounts for the total lateness of this in fraction-days. Decision variable domains are given in (L7). Note that some lines must not be split into multiple loads, while others may be. When a line with multiple cargo burden types is split between multiple asset lifts, each lift is allocated some proportion of the line, and each cargo type is allocated to that lift in that proportion.

Variables *EARLY*, *TARDY*, *LATE*, and *DROP* and the constraints (L1) and (L6) that define them are exhibited only for clarity: In practice, these dependent expressions are eliminated, with their influence added to the objective function (L0), (L2) are relaxed to an inequality (pack) form, and (L6) are relaxed entirely. For clarity of exposition, not all domain constraints on index use are shown, and elastic features for violating other constraints, especially vulnerable throughput capacity constraints (L3), are not shown.

Although we do not need greater resolution for our scenario, there are two embellishments that might appeal elsewhere.

We can replace the $LIFT_{a,d,s,t}$ decision indicating asset *a* departs on day *d* from node *s* bound for node *t* with $CYCLE_{a,d,s,t,s',t'}$ to also keep track of where and when the asset returns for a next load. In our scenario, there is scant difference in transit times from any source node to any destination node, so there is no motive to track complete asset cycles to and from deliveries.

We can augment the decision to $LOAD_{l,a,d}$ some fraction of line *l* on asset *a* for departure on day *d* with a binary $ASSIGN_{l,a,d}$ variable indicating that some volume has been loaded. This permits, for instance, limiting the number of assets among which any given line can be split.

There are also alternate versions of the control constraints (L4) that may be stronger in some applications. For instance, each pair of terms in (L4) can be isolated in its own pair-wise packing constraint.

C. MATHEMATICAL FORMULATION OF ATTACKER, AN INTEGER LINEAR PROGRAM TO MAXIMIZE MINIMUM DISRUPTION FOR AN OPTIMIZED TPFDD

In addition to nomenclature of the target ILP LIFTER, we need:

Index Use [~cardinality]

$\partial \in \Delta$	counts days relative to an attack, i.e. $\partial \in \{0,1,2,3,4\}$
cut	iteration (alias γ) (also cardinality of diversity cuts) [~1,000]

Given Data [units]

$target\_importance_{d,t}$	day *d* utilization of destination node *t* by the best TPFDD plan [value-units]
$damage_{d,d+\partial,t}$	fraction of $target\_importance_{d+\partial,t}$ extinguished ∂ days after an attack on day *d* at destination *t*.
$\overline{attacks}$	maximum attacks [cardinality]
$\overline{daily\_attacks}$	maximum attacks in any one day [cardinality]
$\underline{days\_between\_attacks}$	minimum days between any two attacks [days]

Decision Variables [units]

$ATTACK_{d,t}$	=1 if destination *t* attacked on day *d* [binary]

Derived data:

$attack_{d,t}^{cut}$	Solution $ATTACK_{d,t}^{*}$ from prior iteration *cut* [binary]

Formulaton

$$\max_{ATTACK} \sum_{d,t}\left(\sum_{\partial}\left(damage_{d,d+\partial,t}\,target\_importance_{d+\partial,t}\right)\right)ATTACK_{d,t} \qquad \text{(A0)}$$

$$s.t. \quad \sum_{d,t} ATTACK_{d,t} \le \overline{attacks} \qquad \text{(A1)}$$

$$\sum_{t} ATTACK_{d,t} \le \overline{daily\_attacks} \qquad \forall d \qquad \text{(A2)}$$

$$\sum_{d \le \delta \le d+\underline{days\_between\_attacks}-1,t} ATTACK_{\delta,t} \le 1 \qquad \forall d \qquad \text{(A3)}$$

$$\sum_{\substack{d,t \\ |attack_{d,t}^{\gamma}=0}} ATTACK_{d,t} + \sum_{\substack{d,t \\ |attack_{d,t}^{\gamma}=1}} (1-ATTACK_{d,t}) \ge 1 \qquad \forall \gamma < cut \qquad \text{(A4)}$$

$$ATTACK_{d,t} \in \{0,1\} \qquad \forall d,t \qquad \text{(A5)}$$

1. ATTACKER Formulation Discussion

Given some minimum-penalty TPFDD plan from LIFTER as a target, ATTACKER finds a maximal-penalty disruption of that TPFDD plan.

LIFTER and ATTACKER are applied alternately, with each LIFTER solution seeking a TPFDD plan revision that overcomes the latest proposed ATTACKER disruptions, and each ATTACKER plan seeking some new way to disrupt the best TPFDD plan.

ATTACKER objective (A0) evaluates the damage induced by any attack plan. (A1) limits the total number of attacks. A constraint (A2) limits the attacks in any one day. A constraint (A3) limits the number of days after an attack before another might be mounted. A constraint (A4) is accumulated from each prior ATTACKER plan and forces the current plan to differ in at least one detail from that prior plan.

Although in principle any single attack can inflict simultaneous damage on more than one vulnerable destination node that lasts an arbitrarily long time after the attack, for exposition here we just use $damage_{\partial}$, for $\partial \in \{0,1,...,4\}$ (i.e., we simplify to one kind of attack that has the same consequences on throughput no matter when or where it is mounted.)

THIS PAGE INTENTIONALLY LEFT BLANK

III. COMPUTATIONAL RESULTS

A. JFAST RESULTS (NO INTERDICTION)

JFAST delivers all air requirements to designated debarkation airports within delivery time windows and without shortfall, or dropped lines. As long as cargo arrives in theater within its arrival delivery time window. JFAST reports it “on time.” JFAST fails to deliver all sea requirements. No sea requirements arrive in theater late, but there is a transportation shortfall (that we call “drop” here) of 7% of the general breakbulk cargo, or 553 mTons.

1. JFAST Air plan

JFAST schedules 247 air missions to satisfy 100% of air requirements (see Table 7). Passengers arrive in theater on days C+25 and C+28, while cargo arrives steadily from C+34 through C+37 (see Figure 10). We initially use the full air scheduler to plan the deployment, however because there are no en-route airfields specified in our training TPFDD, JFAST defaults to the quick scheduler. Airfield constraints for passengers and cargo are activated, thus limiting passenger and cargo throughput to planner-specified capacities. *The objective of the quick scheduler, like the full air scheduler, is to minimize overall mission time with respect to asset hours.* With the quick scheduler, each aircraft has a load capacity remaining after deducting fuel weight for the critical (longest) leg of a mission, as well as by actual cabin capacity [Jones 2005]. However, because the quick scheduler uses a flying time estimator, available cabin load seen by JFAST is not visible to us.

Determining airport throughput capacities requires coordination between operational planners and Air Mobility Command (AMC). Airport capacities here are contrived using New York’s John F. Kennedy airport as an example with data acquired from the Port Authority of NY and NJ [Port Authority NY, NJ, 2005] (see Table 8).

	Number of Aircraft Apportioned	Number of Aircraft Used	Number of Missions Flown
C-5	57	4	4
C-17	17	17	22
C-141B	123	78	194
LRWC	30	1	1
LRWP	30	21	26
Totals	257	121	247

Table 7. Aircraft available and used by JFAST. Apportioned aircraft are available to be activated, if needed. No aircraft is allowed to fly more than three missions during planned deployment. E.g., JFAST activates all seventeen C-17, flying a total of 22 missions.

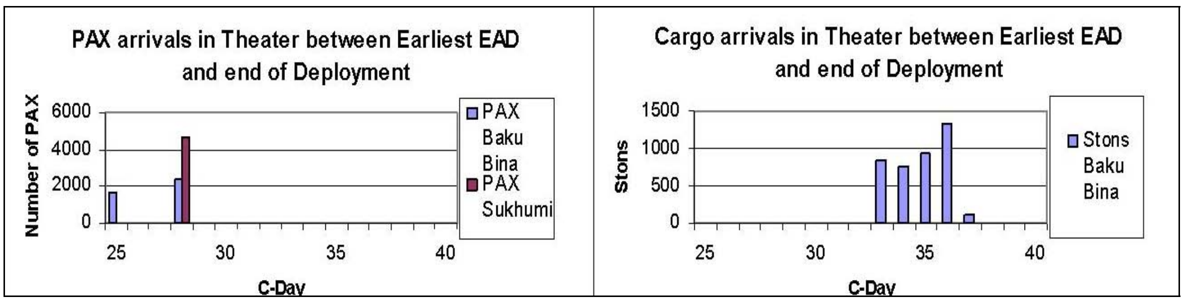

Figure 10. JFAST delivers the passengers and cargo shown to debarkation airports from earliest arrival date through the duration of the deployment. Passengers arrive on two different days, while cargo arrives on five days.

	PAX	***Cargo (sTons)***
Sukhumi Dandra	14,000	*4,911*
Baku-Bina	14,000	*4,911*
Totals	***28,000***	***9,822***

Table 8. Daily airport throughput capacities for passengers and cargo are derived from New York's John F. Kennedy Airport.

2. JFAST Sea plan

JFAST schedules 9 ships, each conducting one mission, to move all but 7% (553 mTons) of the general breakbulk cargo with an average ship capacity utilization of 85% (see Figure 11 and Table 9). Cargo arrives at Poti starting on the earliest arrival date for sea requirements, C+22, and concludes on C+30 (see Figure 12).

JFAST drops cargo lines with no apparent reason. To avoid this, we experiment and force JFAST to mix ammunition with other cargo and to containerize all cargo that can be containerized. Independently, JFAST loads non-homogenous cargo within shipboard compartments, which is forbidden: we did not force this. For example, requirement "TAJFC_R1D" is an ambulance for a medical support group from Fort Bragg. JFAST loads this as measured tons within a breakbulk compartment, rather than in a RORO or LOLO compartment as the TAJFC_R1D cargo category code dictates.

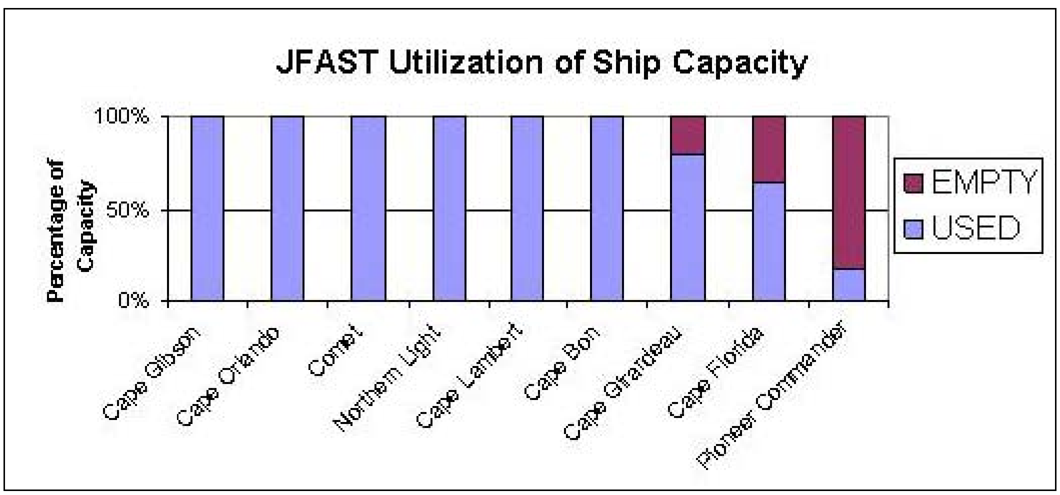

Figure 11 JFAST utilization of shipboard capacity for the deployment. Six of the nine ships activated are fully-loaded. In loading these ships, JFAST has violated some restrictions of cargo categories that should not be commingled in ship compartments.

	JFAST Ship Compartment Utilization	
	Total Capacity	**Percent capacity used**
Breakbulk (mTons)	117667	83.75
Container (MTons)	9847	29.56
Load On Load Off (Sqft)	78928	55.56
Roll On Roll Off (Sqft)	291000	100.00

Table 9. JFAST ship compartment utilization. Some JFAST loads violate compartment cargo class restrictions.

JFAST recognizes a "stow factor" that describes the efficiency with which cargo can be loaded into a compartment. For example, if a cargo compartment has 100 measured tons of capacity, it can hold 75 measured tons of cargo with stow factor 0.75.

We force JFAST to honor a zero tolerance for early arrivals. The minimum load for any vessel in the scenario is set to zero percent of loading capacity. This conveys that that all the cargo in the TPFDD is necessary and we will do whatever it takes to get it to the theater.

Seaport capacity constraints are conventionally provided by Surface Deployment and Distribution Center (SDDC). For this scenario, debarkation seaport throughput capacity is contrived using unclassified data available for Poti Port [Port of Poti, 2005]. The data available for Poti measures throughput capacity in metric tons which differs from the TPFDD that measures sea lines in measured tons, a unit of weight versus a unit of volume. These two varying measures make it difficult to compare JFAST and LIFTER. To reconcile, we do not constrain Poti throughput capacity in JFAST and LIFTER for the non-interdicted run analysis to allow for comparison.

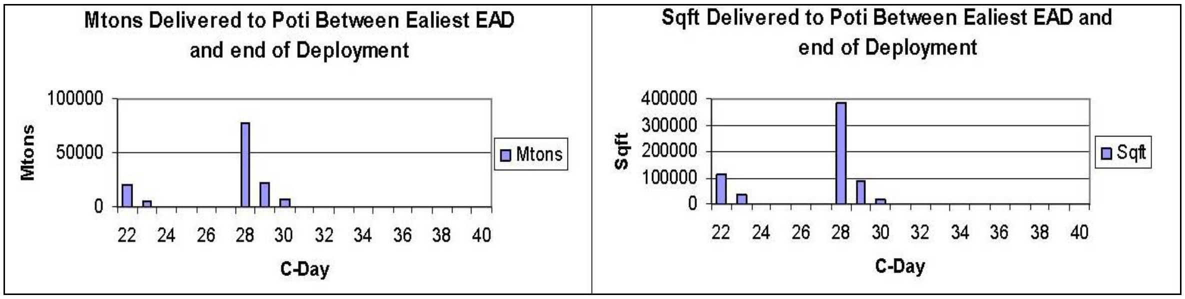

Figure 12. Measured tons and square feet delivered by JFAST to Poti, Georgia, between earliest arrival date and the end of the deployment. All but 7% (553 mTons) of general breakbulk cargo is delivered by JFAST.

B. LIFTER RESULTS (NO INTERDICTION)

We generate LIFTER using GAMS (General Algebraic Modeling System [Brooke, et al. 1998]) and solve it with CPLEX 9.0 [ILOG 2003] on a 2 GHz laptop computer operating under Windows XP [Microsoft 2005]. The GAMS implementation aggressively filters index tuples.

Table 10 shows the penalties we use.

Early penalty	θ_e	1.0
Tardy penalty	θ_t	0.1
Late penalty	θ_l	1.0
Drop penalty	θ_d	10,000
Empty space penalty	θ_u	10.0

Table 10. LIFTER penalty values. There is a large penalty assessed for a dropped line.

LIFTER optimization results resemble those of JFAST. However, LIFTER delivers all air and sea requirements to theater on time. There are no dropped lines and no lines arrive late. Like JFAST, we record when lines arrive and penalize the model for lines that arrive before their arrival time window or after this time window. However, unlike JFAST, we record when within the time window each line arrives. A line that arrives within its delivery time window, but not on the first day of this, is considered tardy. By penalizing tardy lines we give LIFTER an incentive to deliver lines on their

earliest arrival date. Empty space on assets is also penalized. This gives LIFTER an incentive to maximize cargo capacity usage of activated assets.

Although our TPFDD extract does not designate whether a line should be shipped as a unit load (i.e., not split up among assets for partial line shipment), we think such a designation is desirable to restrict optimized solutions from being too clever about filling assets for each lift. Accordingly, we find what fraction of compatible asset capacity each line would occupy, and if a line would occupy no more than a quarter of the smallest compatible asset, we designate the line as a unit load. For our scenario, this rule produces the results shown in Table 11.

Unit Load and Non-Unit Load Shipment Plan					
LIFTER		**Unit-Load Lines**	**Lines Free to be Split**	**Lines Split**	**Split-Line Shipments**
Disaggregated	**Air: AK**	81	110	47	247
	Sea: SE	244	204	5	11
Aggregated	**Air: AK**	1	14	13	265
	Sea: SE	3	14	3	9

Table 11. Unit load and non-unit load shipments. Unit loads must be shipped in a single asset lift, while non-unit loads may be split among asset lifts. Here we define as a unit load any line that occupies less than 25 percent of the smallest asset that might carry it. Surprisingly, even though 204 SE lines are permitted to be split, only five are split, four into a pair of loads, and one into three loads. The "*level 2.5*" aggregation has fewer unit loads, and is only a *slight* relaxation of the 639 level three lines.

By resorting to extraordinary cargo-handling measures, we may be able to load or unload more than the routine facility capacities, albeit at a high cost, otherwise we may choose to make early or late delivery, or drop a load entirely. LIFTER features a mechanism for each of these alternatives, and we specify a penalty per off-load unit violation and penalties for early or late delivery, or dropping a load, and let LIFTER choose the least painful alternative. We could also model secondary debarkation points and alternate routing of cargo, but our focus here is strategic lift and its vulnerability: We want to discover just how good JFAST deployment plans are, and whether they can be made better with or without insurgent threats.

Our specimen TPFDD consists of 639 line requirements (191 air and 448 sea). We solve this seminal problem, but for convenience we also pursue a 32-line aggregate

as mentioned in chapter two. This aggregation necessarily relaxes our "unit load" refinement, which we introduce merely to avoid too-clever optimization that might distribute a line over many assets. We utilize this aggregated TPFDD for interdiction analysis because it is so trivial to optimize as a integer linear program, but we calibrate our aggregate results with the full 639-line specimen.

1. LIFTER Airlift Optimization

We have conservatively assumed that the cargo capacity of each aircraft is reduced by a maximum fuel load at takeoff. LIFTER computes the available cabin load by subtracting the sum of aircraft operational weight and maximum fuel weight from the maximum gross takeoff weight. Load class capacities for bulk, over-size and out-size cargo are reduced proportionate to the reduced available cabin load capacity.

LIFTER airlift optimization delivers 191 lines by activating 119 aircraft from its allocation of 257 and scheduling 261 air missions (compare Table 12 with JFAST results shown in Table 7). (Alternately, we can designate active aircraft to verify JFAST planning results, but LIFTER's suggested plan is so close to JFAST that this is hardly worth the effort, which would require reverse-engineering JFAST fuel management.)

LIFTER satisfies 100% of air requirements. Passengers arrive on days C+25, C+28 and C+29 while cargo arrives every other day from C+33 through C+38 (compare Figure 13 with Figure 10).

The LIFTER airlift ILP has 4,000 constraints, 78,000 variables, 20,000 of which are binary and generates and solves in about 10 seconds. Conversely the LIFTER airlift ILP with the disaggregated TPFDD has 4,200 constraints, 596,000 variables, 288,000 of which are binary, and generates and solves in about 30 seconds.

Most important, *LIFTER gives us a quality guarantee for our plan*. Given the data at hand, and all the assumptions explicitly stated in our formulation, we know with mathematical certainty that there is no better plan as yet undiscovered. To the degree that our plan resembles that of JFAST, we gain some comfort about the JFAST heuristic plan. Given that one of the most important uses of JFAST is to assess lift requirements, and

that JFAST offers absolutely no assurance that its plan is as good as can be, this quality guarantee is a new and important way to calibrate and evaluate JFAST.

	Number of Aircraft Apportioned	Number of Aircraft Used	Number of Missions Flown
C-5	57	3	8
C-17	17	17	39
C-141B	123	77	186
LRWC	30	1	2
LRWP	30	21	26
Totals	257	119	261

Table 12. LIFTER aircraft utilization. Like JFAST, LIFTER uses all 17 C-17's, but flies more missions with each of them. LIFTER favors the larger C-5 and C-17 aircraft more than JFAST. We can control this, and even force LIFTER to use the same aircraft as JFAST.

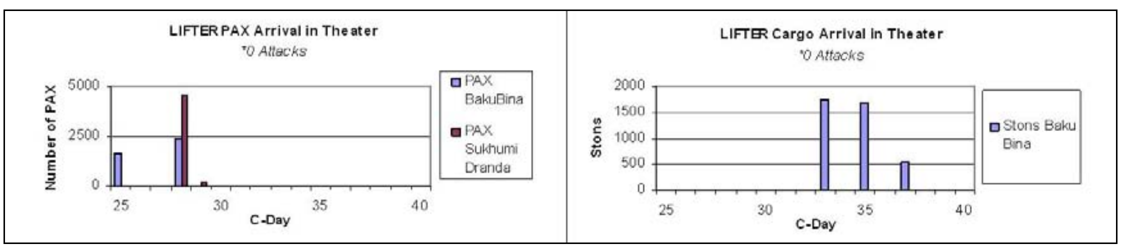

Figure 13. LIFTER passenger and cargo delivery to debarkation airports from earliest arrival date through the duration of the deployment. All passengers arrive in theater on two different days, while cargo arrives in theater on three days, every other day for a span of 5 days.

2. LIFTER Sealift Optimization

LIFTER sealift optimization activates nine ships from its allocation of 91 and breaks 448 lines into 260 loads to satisfy 100% of the sea requirements (compare Figure 14 with Figure 11, and Table 13 with Table 9) with an average ship utilization of 57% of capacity. (Alternately, we can designate active ships to verify JFAST planning results, but LIFTER's suggested plan is so close to JFAST that this is hardly worth the effort, which would require reverse-engineering JFAST's violations of cargo category loading restrictions into ship compartments.)

All cargo is delivered on time and arrives as shown in Figure 15 (compare with Figure 12). LIFTER delivers requirements on days C+22, C+28 and C+29, while JFAST delivers a steady stream of cargo on days C+28, C+29 and C+30. Throughput capacity at Poti is unrestricted in the un-interdicted case to allow for comparison with JFAST.

The LIFTER sealift ILP has 870 constraints, 80,000 variables, 22,000 of which are binary, and generates and solves in about 30 seconds. Conversely the LIFTER sealift ILP with the disaggregated TPFDD has 1,300 constraints, 159,000 variables, 83,000 of which are binary, and generates and solves in about 2 minutes.

LIFTER honors a strict loading policy that dictates how each cargo category must be loaded into compatible ship compartments (see Figure 16). This signal restriction accounts for what superficially appears to be less than optimal ship utilization by LIFTER. This plan is optimal, and strictly adheres to all cargo category load restrictions in ship compartments.

In reality, cargo designated for the lightly-loaded ships might be transferred to other ships by scrupulous violation of load restrictions: An experienced load master can beat a planning model every time by violating such rules.

Here, each ship must load at exactly one seaport: in reality, we might have ships make more than one loading port call, though this is not preferred practice.

Paradoxically, this blemish of lightly-loaded ships is an artifact of the small size of our TPFDD. Larger deployments offer more degrees of freedom to fill activated ships.

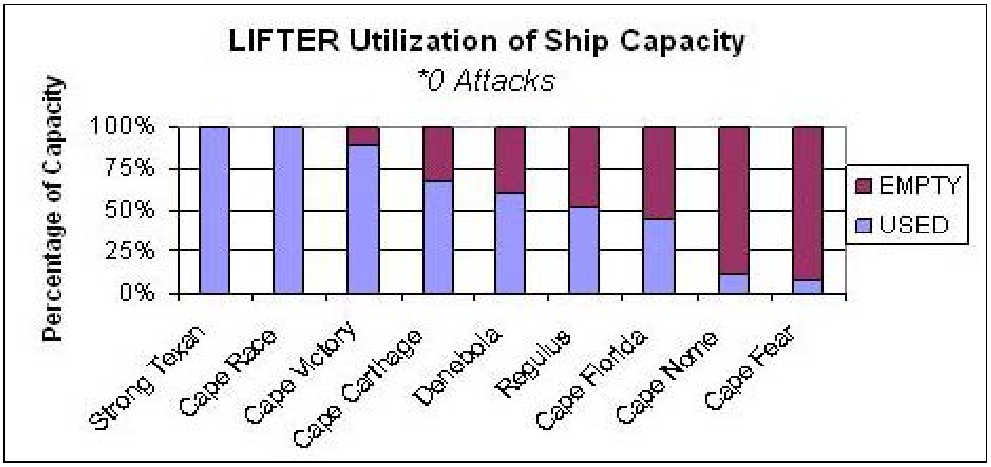

Figure 14. LIFTER sealift optimization utilization of shipboard capacity. Two of the nine ships activated from the 91 ships allocated are completely full. This plan is optimal, and strictly adheres to all cargo category load restrictions in ship compartments. In reality, cargo designated for the lightly-loaded ships might be transferred to other ships by scrupulous violation of load restrictions: An experienced load master can beat an optimization model every time by violating such rules. Here, each ship must load at exactly one seaport: in reality, we might have ships make more than one loading port call, though this is not preferred practice. Paradoxically, this blemish of lightly-loaded ships is an artifact of the small size of our TPFDD. Larger deployments offer more degrees of freedom to fill activated ships.

	LIFTER Ship Compartment Utilization	
	Total Capacity	**Percent capacity used**
Breakbulk (mTons)	50,604	24.0
Container (mTons)	43,969	23.0
Load On Load Off (Sqft)	661,970	91.0
Roll On Roll Off (Sqft)	62,301	42.0

Table 13. LIFTER sealift optimization compartment utilization. LIFTER strictly honors restrictions on which cargo categories can be loaded into each compartment type.

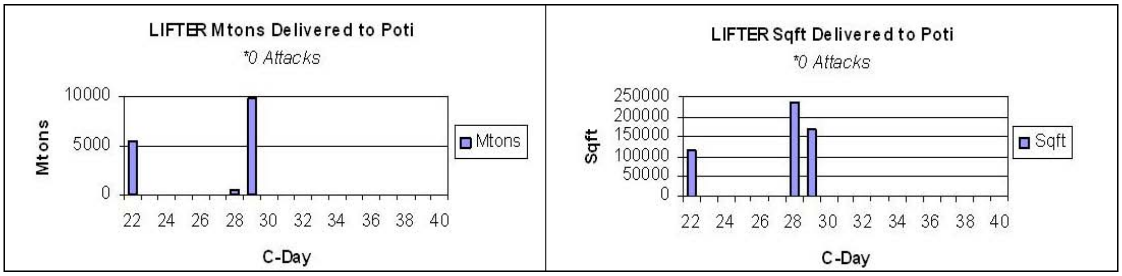

Figure 15. LIFTER uninterdicted delivery profile. Measured tons and square feet of cargo delivered to Poti, Georgia, between earliest arrival date and the end of the deployment. 100% of sea cargo arrives in theater, on time.

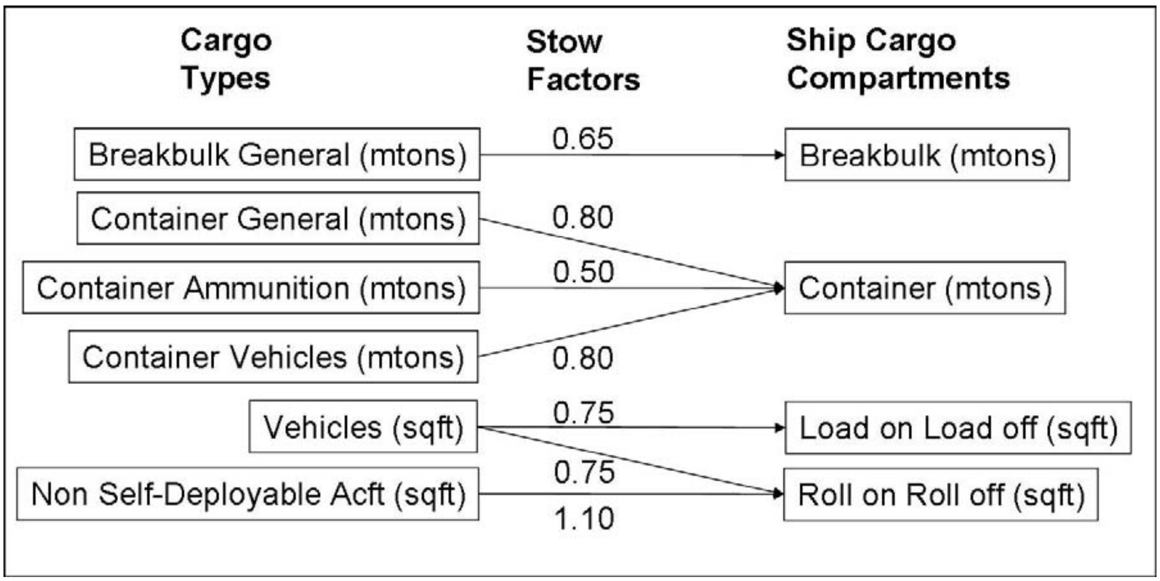

Figure 16. LIFTER strict ship loading policy. LIFTER restricts cargo types that can be loaded into ship cargo compartments. For example, one measured ton of container ammunition occupies twice as much container storage capacity on a ship. This reflects the manner in which ammunition must be safely handled and stowed.

C. LIFTER RESULTS WITH ATTACKER (OPTIMAL INTERDICTION)

ATTACKER achieves its best three-attack interdiction by targeting Poti on days C+20 and C+27, and Baku Bina on day C+35. This optimal attack forces a small amount of air cargo to be dropped: LIFTER revises air missions, but can't avoid dropping 1% of passengers (53), and 1%, or 27.14 sTons, of materiel. Despite the throughput damage inflicted at the Poti seaport, LIFTER revises deliveries to satisfy 100% of the sea requirements anyway.

1. LIFTER Airlift Optimization Responding to the Best Three-Attack Interdiction

Here we have a case where extraordinary cargo handling measures come into play: As part of the optimal attack, ATTACKER completely extinguishes the nominal throughput capacity of Baku Bina on day C+35 and two days hence. But, Figure 17 shows cargo still arriving in theater on these days. In this case the less painful course of action is to pay a high surcharge for exceeding throughput capacity (this is perhaps interpretable as a surrogate for using an alternate debarkation node at higher cost), rather than the even higher penalty of dropping these lines wholesale. We control this decision with our choice of throughput penalties and drop penalties.

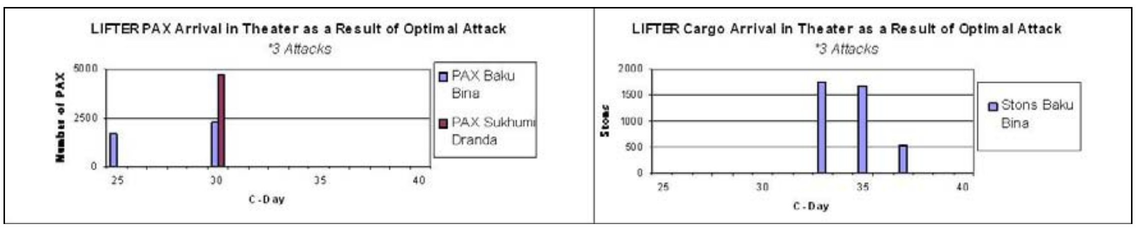

Figure 17. Air cargo delivery profile suggested by LIFTER, responding to an attack on Baku Bina on day C+35. LIFTER chooses to pay the per-sTon penalty of 20,000 units instead of dropping the line for a one-time penalty of 10,000.

ATTACKER iterates over distinct attack plans, tracking the revisions of LIFTER, and discovering its best attack plan after 47 iterations. We see that the delivery profile remains very similar to the un-interdicted case (compare Figure 17 with Figure 13), albeit with much higher penalties

2. LIFTER Sealift Optimization Responding to the Best Three-Attack Interdiction

LIFTER again schedules nine ships with a slightly higher average ship capacity utilization of 58% (see Figure 18) to satisfy the sealift requirements.

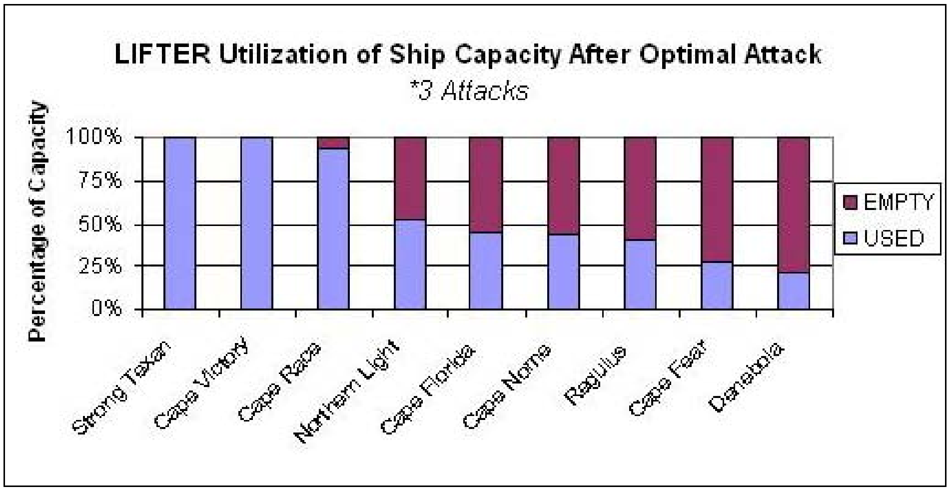

Figure 18. LIFTER ship capacity utilization in response to an optimal three-attack plan. Ship capacity utilization is about the same as in the uninterdicted case.

Given the large capacity at Poti port, LIFTER reschedules cargo delivery to evade ATTACKER disruptions (compare Figure 19 to Figure 15). Cargo arrives in theater on days C+23 and steadily on days C+30 through C+32. This profile completely avoids reduced throughput inflicted by ATTACKER.

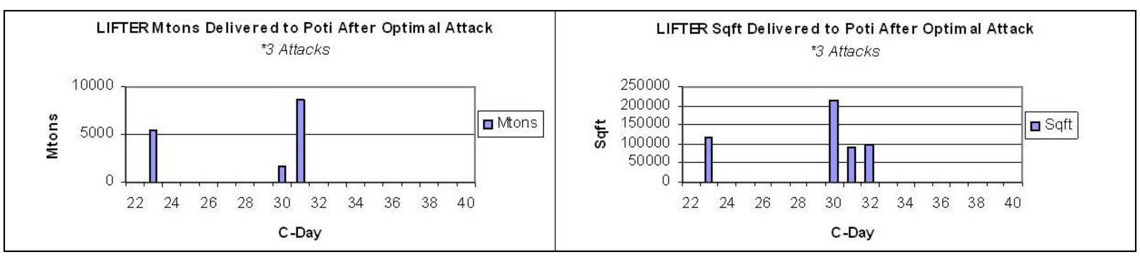

Figure 19. LIFTER sea cargo revised delivery profile responding to the best three-attack plan. Throughput capacity is restricted during days C+20 through C+24 and during days C+27 through C+31. Given the large throughput capacity of this seaport, LIFTER is able to circumvent adverse effects.

D. VARYING ATTACKER CAPABILITY

We analyze one- and two-attack interdictions, realizing that an insurgent may not be able to mount three attacks over the deployment horizon. The results are similar for these more-restrained attacks: When ATTACKER chooses an attack on Baku Bina airport, LIFTER suggests suffering penalties for capacity throughput violation rather than

dropping lines. So, these revised profiles do not vary much, although their costs increase as pushing cargo through damaged facilities gets harder. Our penalties suggest re-routing these aircraft to more-expensive in-theater "pseudo-airfields," rather than drop their lines.

When ATTACKER targets the seaport Poti, LIFTER is able to evade either the one- or two-attack scenarios (see Figure 20). LIFTER reschedules delivery dates to work around diminished throughput capacity and delivers all requirements for the TPFDD.

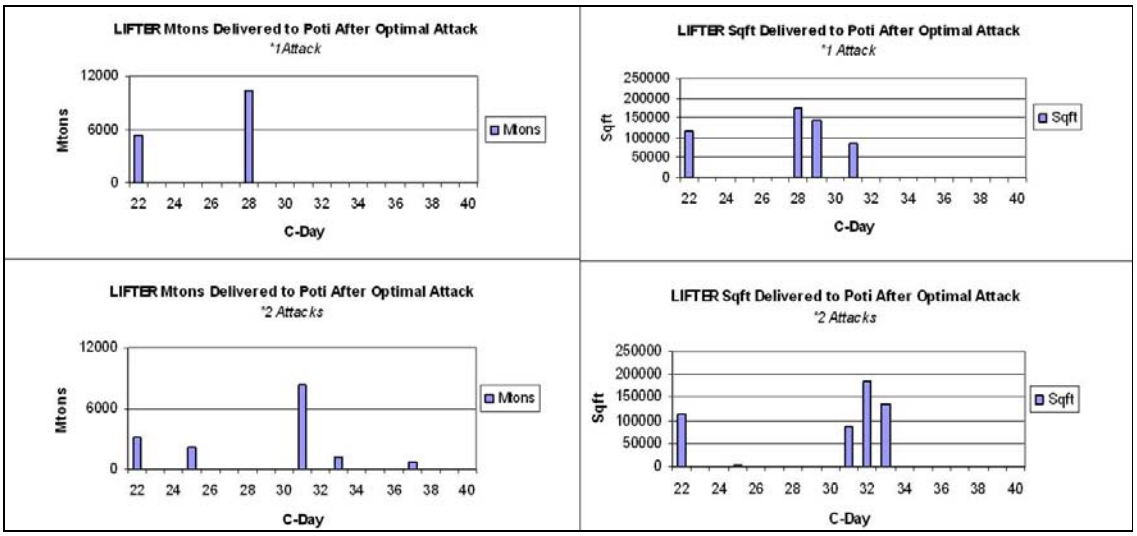

Figure 20. LIFTER sealift delivery profile revisions in response to one attack on C+35 and two attacks on C+28 and C+35. LIFTER shuffles delivery dates to satisfy 100% of sea requirements.

E. NON-AGGREGATED VERSUS AGGREGATED TPFDD LINES

Interdiction analysis using the original non-aggregated TPFDD produces similar results, even though different attack strategies were formulated by ATTACKER (see Table 13).

Optimal Attack Plans for Disaggregated and Aggregated TPFDD lines				
	Disaggregated (639 Lines)		**Aggregated (32 Lines)**	
1 Attack	**Baku Bina**	**C+35**	**Baku Bina**	**C+35**
2 Attacks	**Poti**	**C+27**	**Poti**	**C+28**
	Poti	**C+33**	**Baku Bina**	**C+35**
3 Attacks	**Poti**	**C+22**	**Poti**	**C+20**
	Poti	**C+27**	**Poti**	**C+27**
	Poti	**C+32**	**Baku Bina**	**C+35**

Table 14. Optimal attack strategies derived from the original, disaggregated TPFDD lines, and the nearly-equivalent aggregation of these. The attack strategies differ in detail, but not in effect.

Using either the disaggregated or aggregated TPFDD deployment, ATTACKER chooses to target Baku Bina on C+35 as the optimal one-attack plan. The results of this attack with respect to the responding, reoptimized delivery profile, assessed penalties, extraordinary cargo handling means, and overall cost of deployment are nearly exact. This is reassuring evidence that the aggregated TPFDD produces useful results.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. CONCLUSIONS AND INSIGHTS

A. SUMMARY

We seek an optimal deployment plan that minimizes penalties for early, tardy, late, or non-delivery of personnel and materiel to the theater --- *a truly optimal plan with a guarantee that we haven't overlooked some better one*.

Mimicking current planning practice, we optimally activate assets, load TPFDD lines, and plan deliveries. This is a "secret deployment" that completely ignores any insurgent threat.

We then seek an optimal "surprise insurgent attack" on our optimal "secret deployment," assessing our vulnerability to such a surprise: We allow the insurgents to anticipate our deployment, and schedule attacks on our debarkation points to deny us throughput capacity and maximize the same penalties we have planned to minimize.

Next, we evaluate a transparent, two-sided conflict where the insurgent attacks our deployment and we optimally re-plan around the reduced debarkation throughput capacity denied us. The attacker tracks our revisions and modifies his insurgency. And so forth.

Our goal is to gain insight into deployment planning in a world where "secrets" are hard to keep, and a military deployment is almost surely going to be signaled or detected long before the bulk of our military forces arrive in theater.

A captured Al Qaeda training manual (Department of Justice, 2004) advises: "Using public sources openly and without resorting to illegal means, it is possible to gather at least 80% of information about the enemy." We acknowledge this might be true, and if it is we want to be able to plan accordingly.

B. JFAST

JFAST is evidently the current planning tool of choice for evaluating TPFDD deployments, and so we have adopted JFAST as our reference standard for evaluating any new contribution we suggest.

We have received no formal training about JFAST, relying instead on all the documentation we can muster.

As we have used JFAST, we have forwarded questions to a JFAST training point of contact at USTRANSCOM, who has graciously responded to our many questions with advice and undocumented detail.

We read documentation carefully, and have the unusual advantage gained by mathematical modeling --- if we capture the essence of what JFAST says it's doing, we should be able to mimic similar behavior.

Overall, our research has evidently substantially reverse-engineered the essence of the strategic deployment JFAST function. However, there are some niggling, essential details we cannot recover.

We had a hard time recovering, e.g., air load manifests, by the only means we have been told this can be done: "animate the deployment, slow things down, and click on the aircraft icons." This is tedious for hundreds of missions.

The rainbow chart in the sea model is marginally more convenient. Through either visualization playback, or the sealift rainbow chart, lines loaded on specific assets can be reviewed in the JFAST TPFDD editor ONLY via mouse-clicks. However, the viewable lines are TPFDD summary level two reporting. We need to see level three lines loaded on assets. We also need to export this information directly to, e.g., the tools in Microsoft Office [Microsoft Corporation, 2005] for further analysis.

We develop a military "importance" for each line in the TPFDD. This importance is what LIFTER and ATTACKER both share as the military value of a line reaching the theater, early, exactly on time, tardy, late, or not at all. We cannot recover how JFAST prioritizes the cargo it schedules. We do have access to the line attributes JFAST uses to prioritize lines, but the JFAST evaluation method remains a mystery.

Evidently, JFAST uses a greedy heuristic for scheduling asset cycles. We conclude that this is the case after much experimentation with JFAST, but there is evidently no documentation at all about how this scheduling methods works. We are not the only researchers who would like to understand how JFAST works so that we might evaluate results or suggest improvements: similar findings are reported by

(internationally-renowned, extremely experienced, influential researchers) Morton, Wood and Salmeron [2002] and McKinzie and Barnes [2004].

We have observed a number of cases where JFAST unexpectedly and spectacularly fails. For instance, JFAST has ignored available ships while over-using others. JFAST "rule-based ship selection criteria" (sic) evidently penalizes ship activation in some fashion [Jones, 2005]. JFAST continues using an activated ship to make multiple cycles and deliver materiel very late, rather than activate another available ship to deliver on time.

These obvious failures cause us to wonder how many subtle failures have slipped by.

Given that JFAST is used to certify the feasibility of a TPFDD with allocated assets (or even to justify the need for more assets), this seems to us to be a striking deficiency.

JFAST use of stow factors [Jones, 2005] refers to a table that is not available to all JFAST planners. Even if this table is not "available," this information should at least be visible for validation of ship load plans. Lacking this, we are at a loss to explain why JFAST violates what we understand to be rules about commingling cargo.

C. LIFTER AND ATTACKER

LIFTER suggests optimal deployment plans, with explicit, complete documentation of all assumptions, controls, and details. Such documentation makes it possible to tune and influence LIFTER plans by objective analysis. For instance, we can carefully trade off dropping a line with making extraordinary efforts to deliver it, despite interdicted throughput capacity. In the un-interdicted case, LIFTER plans our test deployment with scant penalties, for lines delivered a little tardy, and assets used with empty space. Increasingly aggressive interdictions clearly force LIFTER to make hard decisions.

D. ADVANTAGES OF OPTIMIZATION, VICE HEURISTICS.

A distinguishing contribution of this work is that we offer a truly optimal, face-valid deployment plan, and we also provide an objective assessment of its quality. That is, we clearly state all our assumptions, and then mathematically guarantee that no other plan exists that satisfies these and is better than the one we offer. This has much to recommend it, as you immediately discover when JFAST beautifully animates its planning solutions produced by its undocumented heuristics that without warning exhibit breathtaking errors and oversights. You hope for such obvious telltales of trouble, because there is no other warning that a poor-quality plan has been offered, even when much better plans exist.

E. THE VALUE OF INTELLIGENCE

What is the value of secrecy? If we can conduct our TPFDD lifts in complete secrecy, we can presumably avoid the hazards of insurgents. In our case, this would mean a secret lift incurring a few tardy penalties totaling about 1.5 million "importance-day" units (i.e., our composite measure of the importance of each line multiplied by the days of tardiness). In this case, our only constraints are imposed by asset and throughput availability.

What is the value of complete surprise for the interdictor? If insurgents can interdict our shipments optimally and by complete surprise, they can inflict penalties totaling 520.6 million (see Figure 21).

The transparent case yield penalties totaling 47 million. Clearly, the ability to hide activities can have orders-of-magnitude impact for both sides in this exchange. The transparent result is the best, worst-case outcome, with both sides endowed with omniscience – full visibility of the conflict. These extreme points (see figure 21) --- total secrecy for the deployment, full transparency between opponents, or total surprise for an interdiction --- can be used for excursions representing partial secrecy, or partial surprise. Whether you pursue these nuances, or not, Figure 21 should give pause to anyone contemplating planning a military action assuming complete surprise.

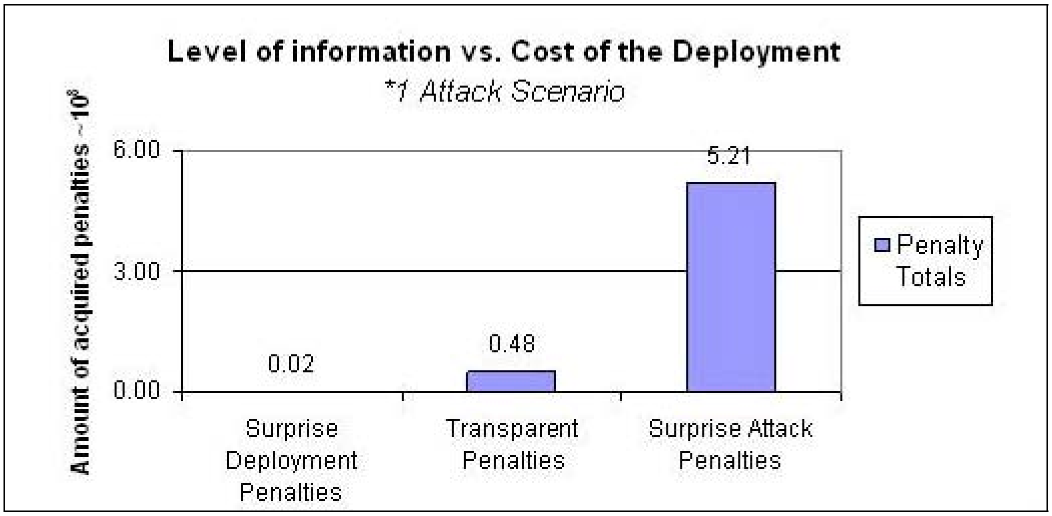

Figure 21. The value of secrecy and intelligence. From left to right, we see the cost of the deployment increases with the amount of information an interdictor possesses. Assuming the interdictor has no information and we surprise him with our deployment, there are negligible deployment penalties (1.5 million). Conversely, if the interdictor has perfect information, and is able to mount a surprise attack against us, the deployment penalties are exceedingly, even distressingly high (521 million). The "transparent" result is what we achieve by planning in anticipation of an interdiction (48 million). The optimistic case at the left is what we assume by ignoring vulnerability to interdiction, the pessimistic case at the right is what might result if this assumption is false.

F. EMBELLISHMENTS AND GENERALIZATIONS

It is easy to extend the geographic span of planned shipments beyond the port of debarkation into the theater, whether such shipment planning preserves the fidelity of discrete lines, or merely represents shipments as gross materiel flows. Brown [1999] shows how to do this, and also shows how to re-route shipments when debarkation capacity is diminished for any period of time. The overland logistics necessary in Georgia and Azerbaijan (see Figure 7) would make such diversions daunting. So, although in theory we could model such evasion of insurgent interdictions, in practice we have chosen to merely re-schedule around the disruptions at the planned debarkation points.

We have shown how to optimize movement of half a division --- 639 level-three TPFDD lines. Do our methods scale up? The dimensions reported here are for the full

allocation of ships and aircraft, and we already include these. The sea and air models are de-coupled. The size of each grows linearly with the number of lines. We are comfortable scaling this up to emergent deployment of several divisions (i.e., more personnel and materiel, but about the same planning horizon). To accommodate even larger, or longer deployment plans, we would abandon the algebraic modeling language GAMS and substitute a direct model generator. In lieu of direct solution of a monolithic problem for sea and air, we would employ indirect means (i.e., column generation). In our experience, this effectively removes size limits.

JFAST is a useful tool for "surprise" deployment planning, albeit dependent upon local search heuristics that suggest plans of unknown quality. JFAST is fast enough to be used repeatedly in a heuristic decomposition that mimics ours. Our optimization model ATTACKER emulates resource-limited interdiction actions, and this model can be heuristically solved with local search. An essential feature of this attack heuristic that is unusual for local search is a filter that guarantees some feature of each successive attack differs from all prior ones. Lacking such a filter, random attacks can be mounted, but this is pretty inefficient. A local, greedy objective suffices in our experiments to deliver good-quality attacks in a few iterations pursuing unique attacks. Each attack, and response, offer a candidate "transparent" plan that is tested to see if a new overall best incumbent has been discovered.

LIST OF REFERENCES

Allard, K., 1995, “Somalia Operations: Lessons Learned,” National defense University Press, Washington D.C..

Bennett, D., 1997, “Defending Inlandia: A flawed Deployment Doctrine,” Joint Forces Quarterly Autumn Winter 1997-98.

Brooke, A., D. Kendrick, A. Meeraus and R. Raman [1998] GAMS: A User’s Guide, GAMS Development Corporation, Washington, DC, September.

Brown, W., 1999, “Re-optimization of Time Phased Force Deployment Plans in Response to Emergent Changes During Deployment,” M.S. Thesis in Operations Research, Naval Postgraduate School, Monterey, CA, September.

Clarke, R., 2005, “Ten Years Later,” The Atlantic Monthly, January-February.

Department of Justice, 2004, “Al Qaeda Training Manual”, http://www.usdoj.gov/ag/trainingmanual.htm , accessed 2 December 2004.

ILOG [2003] ILOG CPLEX 9.0, ILOG S.A., Gentilly Cedex, France.

Jones, Kenneth, 2005, “JFAST 8.0 Training Manual”, http://www.jfast.org/jfast/Manuals/jfast8/WebPageChapters%201-3.htm , accessed 05 March 2005.

Joint Forces Staff College Pub 1, 2000, “Joint Staff Officer’s Guide,” National Defense University Press, Washington D.C..

John Hopkins Applied Physics Laboratory, 2005, “JHU/APL M&S Applications,” http://www.jhuapl.edu/programs/modsim/application.html#WLTAE , accessed 05 March.

McKinzie, K., Barnes, J.W., 2004, “A Review of Strategic Mobility Models Supporting the Defense Transportation System,” Graduate Program in OR/IE Research Paper, The University of Texas at Austin, Austin, TX, July.

Microsoft Corporation, 2005, “Microsoft Corporation,” http://www.microsoft.com/, accessed 17 March .

Morton, D., Wood, K., and Salmeron, J., 2002, “A Stochastic Program for Optimizing Military Sealift Subject to Attack,” Graduate Program in Operations Research Paper, The University if Texas at Austin.

Port Authority of New York and New Jersey, 2005, “JFK Aviation Traffic Reports,” http://www.panynj.gov/aviation/traffic/coverfram.HTM , accessed 05 March.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California

3. Professor Gerald Brown, Code Or/Bw
   Department of Operations Research
   Naval Postgraduate School
   Monterey, CA

4. CDR Gelnn J. Lintz, Code Or/Bw
   Department of Operations Research
   Naval Postgraduate School
   Monterey, CA

5. Professor Matthew Carlyle, Code Or/Bw
   Department of Operations Research
   Naval Postgraduate School
   Monterey, CA

6. Mr. Curtis Badgett
   Joint Warfare Analysis Center
   Dahlgren, Virginia

7. Mr. Lyn Thomas
   Joint Warfare Analysis Center
   Dahlgren, Virginia

8. Diretor of Naval Intelligence (N2)
   DNI Executive Assistant (N2A)
   Washington, D.C.

9. CAPT Alex Butterfield
   COMPACFLT N2
   Honolulu, HI

10. CAPT Stephen Ashby
    Naval Postgraduate School
    Monterey, CA

11. LT Peter M. Koprowski
COMPACFLT N2
Honolulu, HI